Intel Developer Update is Intel's monthly online news magazine for developers. As the official publication of developer.intel.com, it brings hardware, software, and Web developers the latest information on Intel initiatives, technologies, platforms, and products.

Cover Story

Each month, we run a cover story on the most significant industry announcement, trend, or development for the month.

Featured Articles

Delivering in-depth reports on key platforms, products and technologies, our featured articles provide a monthly source of information on issues affecting developers. Be sure to check in every month for the latest developments driving the evolution of the industry.

Contact the Editor

To make *Intel Developer Update* a better information resource, we invite you to share your thoughts on what we've published or what you'd like to see covered. Comments are always welcome.

Archives

Our archives contain two groups of previously published articles. One group contains all the articles that appeared in *Platform Solutions News*, the earlier version of *Intel Developer Update*. The articles date from September 1997 through August 1999. The other group is set up to contain *Intel Developer Update* articles dating from the inaugural September/October 1999 issue.

Bookmarking

We advise against bookmarking article pages. They're accessible online only during the month the issue is live. Thereafter, they're removed to our archives. Instead, we suggest that you bookmark the PDF (Adobe* Portable Document Format) file versions of the articles. You'll find buttons for the PDF files labeled "print article" in the right navigation section of each article. A PDF for the entire issue is labeled "print magazine" and is located near top right side of the IDU home page.

Table of Contents

(Click on page number to jump to articles)

COVER STORY

Beyond Gigahertz 3

DEPARTMENTS

DESKTOP

Moving to Socket mPGA478 Desktop Boards 8

INITIATIVES AND TECHNOLOGIES

Intel Press Unveils Four New Technical Books 13

Creating a Third Generation I/O Bus 15

An Integrated Approach to Audio Design 18

New Technology Brings RosettaNet* to Smaller Businesses 21

Creating Peer-to-Peer Middleware from Web Services Technologies 24

NETWORKING & COMMUNICATIONS

Intel® WAN/LAN Access Switch Example Design Highlights the IXA Ecosystem 28

SERVERS

Breakthrough Performance for Intelligent Internet Storage 31

SOFTWARE

The Intel® Integrated Performance Primitives API 35

Automatic Debugging of Multithreaded Programs 38

WEB DEVELOPMENT

Intel Drives Richer Web Services through e-Business Standards 40

WIRELESS

Automated Client-Based Layer-3 Switching 43

Note: Intel does not control the content on other company's Web sites or endorse other companies supplying products or services. Any links that take you off of Intel's Web site are provided for your convenience.

Cover Story

Beyond Gigahertz

Paul Otellini
Executive Vice President
Intel Architecture Group
Intel Corporation

Overview

Intel has unveiled the world's first 2-GHz microprocessor, and at the Intel Developer Forum Conference Fall 2001 we demonstrated an Intel® Pentium® 4 processor running at 3.5 GHz. Over the lifetime of the Pentium 4 processor family, Intel expects this architecture to scale to 10 GHz and beyond.

Intel continues to make substantial investments in product development, manufacturing, and process technologies that will enable us to maintain our leadership in microprocessor speed and performance. While the focus on raw processor speed continues to be important for the industry, it is not the only key metric for the industry going forward. While these gigahertz are necessary, they are not sufficient in themselves to enable rich computing experiences for the next decade.

Intel is focusing its product and technology development vision beyond gigahertz. The goal is to bring fundamental technologies to market that enable additional productivity and enhanced capabilities for computer users.

This focus includes the fundamental technologies and chip design features that will be required to deliver greater value and functionality to the people who use computers every day. Intel's objective is to enable a wide range of usage models with the ability to redefine how people will use the pervasive and powerful computing resources Intel is working to create.

As the computing industry has grown and new technologies have evolved, purchasing criteria are changing. While processor speed is of primary importance, buyers now look to features such as style, form factor, security, power consumption, reliability, communications functions, price, and overall user experience. Combinations of these features and others are driving end user technology requirements in individual market segments. Intel plans to develop technologies that will help address these changing requirements in each of the key market segments.

Intel is also working to accelerate the technologies needed to improve the performance of the computing platform, including bus and interconnect performance. These new technologies include 3GIO, Serial ATA, InfiniBand*, USB 2.0, and AGP 8X. Intel's active role in industry-wide standards groups, along with its own technology development efforts, exemplifies the company's drive to enable an ever-widening range of computing capabilities.

By extending our vision beyond traditional measures of performance and tailoring processor and platform enhancements to specific market segments and customer requirements, Intel will help bring more intelligence, greater capabilities, and a new level of pervasiveness and usefulness to computing.

Hyper-Threading Technology

Two kinds of parallelism exist today to improve processing performance. These include Instruction Level Parallelism (ILP), implemented in super-scalar processors, and Thread Level Parallelism (TLP) implemented in multiple-processor system configurations. While today's processors exploit ILP, mutually exclusive hardware resources may still exist.

Hyper-Threading Technology is the term for an exciting new Intel® processor technology that enables a single physical processor to appear as two logical processors to the operating system and accompanying software. While they share the same physical processor core, the two logical processors can execute different tasks in parallel.

Hyper-Threading Technology increases processor performance by combining the gains achieved through both instruction-level and thread-level parallelism. Virtually all contemporary operating systems divide their workload into processes and threads that can be independently scheduled and dispatched to run on a processor. In addition, Hyper-Threading Technology takes advantage of the thread-level parallelism found in most high-performance applications including database engines, scientific computation programs, engineering-workstation tools, and multimedia programs.

In a hyper-threaded configuration, the threads are dispatched for execution on each of the logical processors. The processor core then executes these two threads concurrently, using its out-of-order instruction scheduler to keep as many of its execution units as possible busy during each clock cycle. One of the principal benefits of Hyper-Threading is that it enables the processor to use on-die resources that may have otherwise been idle.

In enterprise applications, the benefits of Hyper-Threading Technology include improved overall system performance, since tasks can be executed in parallel. The technology also enhances scalability by increasing the number of users that a given physical platform can support as well as by increasing the number of transactions that the platform can execute in a given amount of time.

The processing of multiple threads by a single physical processor represents a new approach to improving the instruction throughput of processors in servers and high-performance workstations. Hyper-Threading provides a glimpse of the future of microprocessor design in which the performance of a processor in executing a specific application may be as important as its clock speed. Like other extensions to the IA-32 microarchitecture, the performance gains resulting from Hyper-Threading are independent of the clock speed of the chip.

The good news for the industry is that software that is multi-threaded for multiprocessor systems today will similarly take advantage of Hyper-Threading-enabled systems. Recompiling the software for Hyper-Threading Technology can deliver additional performance, and vendors can achieve even more significant performance gains by optimizing their software for Hyper-Threading.

Intel is working with third-party vendors to provide Hyper-Threading development vehicles to optimize software development and tune their multi-threaded code for future versions of Intel® NetBurst™ microarchitecture processors with Hyper-Threading Technology.

The first software development vehicles were delivered in December 2000, and they provide software developers with Hyper-Threading-enabled hardware platforms for software testing and tuning. Intel also provides development tools including compilers, a performance analyzer, performance libraries, and Intel Developer Services online.

Software tuning is now underway for software applications including relational/on-disk database, security encryption/decryption, enterprise resource planning and customer resource management, Internet infrastructure, middleware, database in-memory, supply chain management, business intelligence, and voice/computer telephony integration. Intel has also worked with operating system vendors to communicate the benefits of Hyper-Threading.

Developers should start preparing their applications now to take advantage of next-generation Intel processors with Hyper-Threading Technology.

Enterprise Requirements

The increasing data demands in the enterprise segment represent a significant growth opportunity for the industry. The growing pervasiveness of e-Business demands a combination of intelligent compute performance, scalability, and data availability. Over time, Intel is working with the industry to advance server technologies in a number of areas, including:

Performance

- High-performance Intel® Itanium™ processors
- The ability to seamlessly add execution units to increase parallelism
- Hyper-Threading Technology
- Multi-core technology
- High-performance bus interconnect technology

Scalability

- High-bandwidth I/O: InfiniBand architecture, 3GIO components and integrated technologies
- Server clustering technologies
- Support for large memory systems: the Itanium processor family, which can directly address terabytes of memory
- Scale-up SMP systems with 64 processors and beyond
- Availability Extensive Error Correcting Code (ECC) and Machine Check Architecture (MCA) on critical data paths and memories to contain, correct, and recover from data and process errors
- Support for automatic fail-over and hot-swap subsystems
- Advanced thermal management technologies, ranging from next-generation parallel redundant cooling concepts to design considerations including layouts, venting, thermal design power, fans, and ducting
- Integrated architecture reliability, availability, and serviceability (RAS) built into all levels of the platform, including the processor, chipset, firmware, and operating system

The data-handling demands on the enterprise are growing larger every year, with more users, more data, and more data-driven applications. The growth of enterprise technology is enabled by increases in system performance. At the same time, the enterprise information technology infrastructure must be up and running. Enterprise-class servers must be available 24x7x365, or in excess of 60,000 hours per year.

At the IDF Conference, Intel demonstrated some of the RAS capabilities of a computing system based on the next-generation Itanium processor, code-named McKinley, and the Intel® 870 chipset. The demonstration emphasized the reliability features that Intel is building into its high-end server products.

During the demo, a McKinley-based server running a typical enterprise solution stack had intentional errors induced as it processed information. Intel's enhanced Machine Check Architecture (MCA) technology allowed the system to continue executing transactions while it recovered from these error conditions and maintained data integrity.

A key feature of MCA is its ability to detect and correct errors by allowing the process to be recognized by operating system software and other important elements of the server system. MCA is capable of analyzing data and responding to it in a way that enables higher overall system reliability and availability. The Intel® MCA, which is an open technology that developers can customize to meet unique needs, is a key feature that is being built into Intel’s Itanium processor family to help meet the reliability requirements of the systems that businesses and organizations rely upon.

Intel’s open and extensible system architecture enables OEMs to implement customized system reliability features. It delivers best-in-class RAS with the proven economics of Intel® Architecture.

Low Power

As usage models evolve, the need for low power continues to become important across multiple computing segments including ultra-dense servers, blade servers, small form-factor desktops, and mobile PCs.

Of course, low power is not new in the mobile computing segment. The continuing need for more mobility has resulted in thinner/lighter form factors and longer battery life. The mobile form factor has gone through incredible changes over the past five years, from an almost portable machine to a truly mobile PC. Intel expects the need for low power to become even more important as wireless networking technologies become increasingly pervasive on the mobile platform. Thinner/lighter mobile form factors continue to grow in importance as we move to the era of ”anytime anywhere” computing.

In the desktop space, small form factor platforms have become increasingly popular over the last few years. Intel expects this trend to continue going forward. The popularity of small form factor desktop PCs is being driven by the need to save desktop space and the continuing need for greater ease of use. As users demand smaller form factors, the requirement for delivering lower power without unduly sacrificing higher performance becomes more important.

Low power has also been gaining momentum in the server segment over the last 1–2 years, where compute density has become increasingly important. Compute density is a key cost driver for data centers and Intel expects the demand for front-end servers to grow significantly over the next 3–5 years.

So how are Intel and the industry working to meet these requirements for lower power while delivering higher performance?

Intel looks at low power on two vectors:

- First is the need for smaller form factors. Generally, the smaller the form factor, the more difficult it is to cool the processor and the platform, and the greater is the requirement for lower thermal design power. The increasing need for smaller form factors across mobile, desktop, and servers makes this vector important in all segments.
- Second is the need for longer battery life. Battery life is a function of the average power consumed by the processor and the rest of the platform. This is primarily important for mobile PCs, where users are valuing longer battery life for anytime, anywhere computing.

Intel has had a long history of innovation on both of these vectors. As an example, "Voltage Reduction Technology" was first developed in Intel Labs over 10 years ago. This was truly a breakthrough at the time, and over the years this innovation has found its way into our products in a variety of ways. Recent examples include the "Low Voltage Mobile Pentium® III Processor Product Line," which consumes less than 1W average power, and the "Ultra-Low Voltage Mobile Pentium® III Processor Product Line," which consumes less than half a Watt average power.

Clock gating provides another good example. This technology was first pioneered almost seven years ago, and today Intel continues to improve it and deploy it in its processors to turn off the circuits that are not needed to conserve power.

Another example is the Advanced Configuration and Power Interface (ACPI) specification and Intel's involvement with industry leaders such as Microsoft to drive new power-management capabilities into PC platforms.

A good example of Intel's commitment to delivering higher performance at low power is the next-generation mobile processor, code-named Banias, that is being developed in the Israel Development Center. The design team is developing a number of architectural, logic, and circuit enhancements to meet its goal of higher performance at low power.

Examples of these techniques include aggressive clock gating, which turns off parts of the processor that are not being used to save power, special sizing techniques, which enable minimal power consumption in circuits, and micro ops fusion, where instructions are combined for faster execution. Banias is scheduled to be launched in the first half of 2003, with a sustained roadmap over the next several years.

Industry Call to Action

Intel is uniquely positioned to work with the industry as we move into the next age of computing, with competencies in silicon, platforms, validation, and technology initiatives, anchored by wide industry participation.

As we focus on enhancing experiences for computer users in multiple segments, it is important to remember that Intel is not moving away from gigahertz as a key performance element, but rather building on the foundation of processing performance to deliver enhanced computing capabilities across multiple market segments.

It is more important than ever for the industry to work together to build great products to meet enhanced user needs and to meet the requirements of evolving usage models.

The industry can accomplish this by doing four things:

1. Take advantage of the raw MIPS available from Intel processors. Intel will continue to provide leadership in processor technology to empower new levels of application performance for the digital world.
2. Innovate around increasing parallelism by threading software applications, drivers, and operating systems.
3. Utilize new low-power architectures to create breakthrough products.
4. Address the single largest growth opportunity, the enterprise, with the Intel Itanium processor family.

Summary

Intel has launched the world's first 2-GHz microprocessor, and at IDF Fall 2001 we demonstrated an Intel Pentium 4 processor running at more than 3.5 GHz. Over the lifetime of the Pentium 4 processor, this architecture is expected to scale to 10 GHz and beyond.

Intel will continue to invest in the product development, manufacturing, and process technologies that will allow us to lead in microprocessor speed. While the focus on gigahertz has its place, raw processor speed by itself it is not sufficient to drive the levels of growth and innovation that are required for the industry to continue to prosper.

Intel is expanding its product and technology development focus beyond gigahertz to bring to market fundamental technologies that will enable greater productivity and better experiences for computer users. These technologies include parallel processing based on Hyper-Threading Technology, data availability solutions for enterprise-class servers, and new generations of processors that combine high performance with low-power requirements.

Specifically, Intel will create technology capabilities that are targeted across a wide a variety of market segments and computing models. This tailored approach to the development of products and technologies will match the increasingly sophisticated computing requirements of the digital world.

Author Bio

Paul Otellini is responsible for Intel Architecture businesses and strategies. Reporting to him are Intel's computing business units, including the Enterprise Platforms Group, the Desktop Platforms Group, the Mobile Platforms Group, the Technology and Research Labs, IA Strategic Planning, IA Marketing Group, Solutions Enabling Group, Solutions Market Development Group, Microprocessor Marketing and Business Planning, and Reseller Products Group.

Prior to this appointment in January 1998, Otellini was the executive vice president of sales and marketing during which time he focused on extending Intel's global presence into emerging markets and making Intel a leader in the use of e-Commerce.

Otellini joined Intel in 1974 and has held a number of positions during his career. After managing Intel's business with IBM Corporation from 1980 through 1985, Otellini was promoted to general manager of the Peripheral Components Operation and to general manager of the Folsom Microcomputer Division in 1987. Otellini was appointed an operating group vice president in 1988.

In 1989, he became assistant to Andrew S. Grove, then president and now chairman of Intel. Otellini took on responsibility of the company's Microprocessor Products Group in 1990 as general manager, managing the introduction of the Pentium® processor. He was elected a corporate officer of Intel in 1991, was made a senior vice president in 1993, and was promoted to executive vice president in 1996.

Otellini received a B.A. in economics from the University of San Francisco in 1972, and an M.B.A. from the University of California at Berkeley in 1974. He is a native of San Francisco and is actively involved in educational and fine arts initiatives.

Departments

Desktop

Moving to Socket mPGA478 Desktop Boards

Dan Ragland
Technical Marketing Engineer
Intel Architecture Marketing Group
Intel Corporation

Overview

Intel's desktop board roadmap now includes boards that support Socket 478 Intel® Pentium® 4 processors that will scale performance to 2 GHz. These new processors are available in the mPGA478 socket, a smaller form factor intended to be widely adopted as the Intel® product roadmap moves beyond PGA423-based designs.

The good news for system integrators is that for the first time Intel Pentium 4 processor-based desktop boards are ready for integration in microATX form factor systems. These longer life boards are designed to support the Pentium 4 processor roadmap until the fourth quarter of 2002.

For product longevity and stability, OEMs and integrators should convert their Intel® 850 chipset-based systems to mPGA478-based boards as soon as possible. Making a successful transition requires system integrators to understand a new set of design considerations including processor cooling requirements, a new heat sink retention and board stiffening mechanism, microATX power supply specifications, and chassis design guidelines.

The mPGA478-based design features a new thermal/mechanical solution with a new heat sink support mechanism and a chassis-independent attach design that provides integrators with more options than the direct chassis attach mechanism required for earlier PGA423-based platforms.

To help ensure a successful transition to cost-effective and stable microATX PC platforms capable of delivering improved thermal performance and headroom for Intel's Pentium 4 processor roadmap through 2002, it's important to understand the design issues, and then follow with prompt implementation and thorough testing.

Thermal/Mechanical

As shown in Figure 1, a new Thermal Mechanical Solution cools the Pentium 4 processor, and the screws required for the PGA423 package have been eliminated in favor of pushpins.

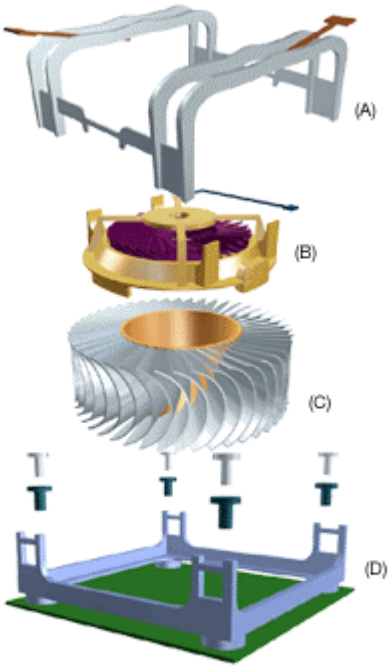

Figure 1. Intel's reference design for a socket 478 thermal mechanical solution. Note: This picture is for illustration purposes and actual dimensions will vary.

In addition to providing extra cooling capacity for the processor, the assembly retains the processor package to the board and provides added stiffening for the motherboard itself. With the elimination of the through-hole mounting technique used in Socket 423-based boards, extra stiffening is now required to protect the surface-mount mPGA478 processor and the chipset and to provide support for the additional weight of the thermal solution.

Information regarding the Intel® Reference Thermal Mechanical Solution is available in the *Performance MicroATX Design Guide 2.0*. Consult your Intel representative for information on where to obtain individual components of the thermal/mechanical solution. Integrators have a variety of implementation options, including solutions from Intel and third-party vendors. The retention mechanism assembly ships with Intel® D850MD and D850MV Desktop Boards, and the processor clip assembly is bundled with retail boxed Pentium 4 processors. Figure 2 depicts the platform assembly flow required for systems with the Thermal Mechanical Solution.

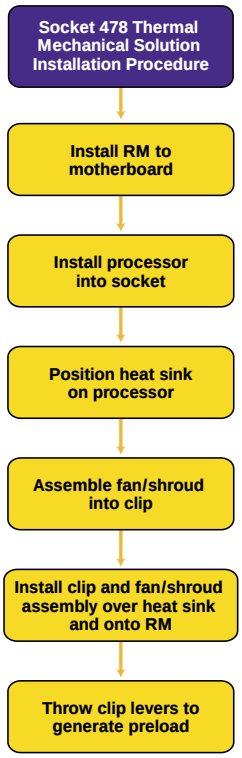

Figure 2. Assembly flow process for systems with the 478 Thermal Mechanical Solution

Figure 3 shows the ATX keep-out specification overlaid on a microATX Board. This layout leads to necessary violations of the ATX specification. Meeting this recommendation requires a chassis with a minimum 1.5 inch clearance over ATX keep-out zones C, D, and E.

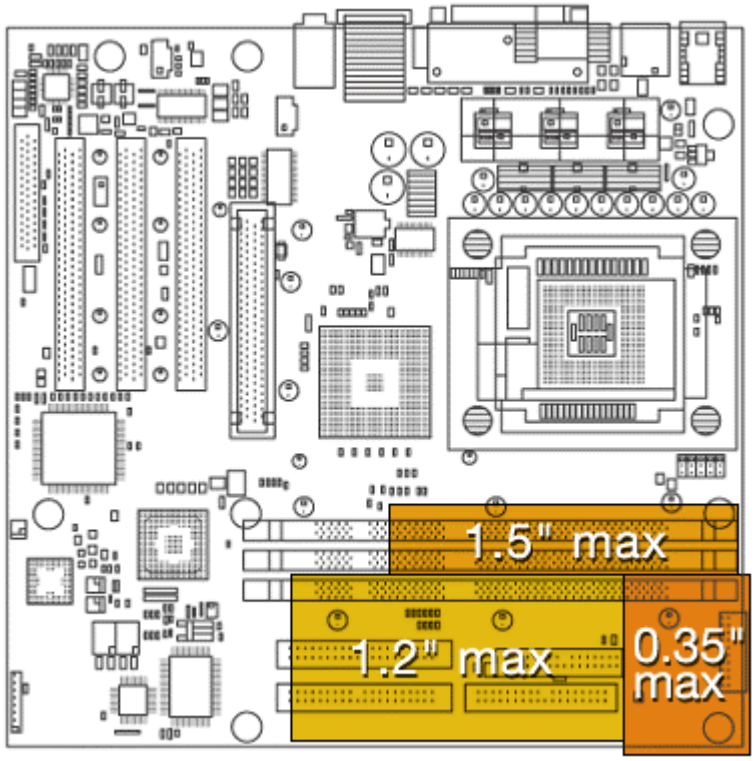

Figure 3. ATX Keep-out zones.

Cooling

Cooling is an important design consideration in systems based on the latest generation of Pentium 4 processors. While thermal issues are more important than ever, meeting these requirements involves familiar methodology.

System builders should refer to the guideline entitled *Performance microATX Thermal Design Suggestions* version 1.0 to understand requirements specific to their configurations.

The guideline specifies a recommended $Tproc_{inlet}$ of 40°C measured one-half inch above the processor intake fan, with temperature readings taken at an ambient room temperature of 35°C.

The chassis should meet the thermal requirements as provided in the *Design Suggestions*, including optimal airflow, supported by a motherboard-powered rear exhaust fan and improved venting. Each completed system design should then be tested to verify that it meets the required parameters.

Intel provides software tools that can help meet system cooling requirements. The Intel® Active Monitor is a software utility that enables system builders and end users to monitor critical system parameters including temperature, fan speed, and power voltages. This utility is included with Intel® Desktop Boards.

Power Supply

Now that Pentium 4 processors can be integrated in microATX platforms, it is important to specify a microATX power supply that meets the requirements of this class of processors while providing power headroom for future growth.

General power supply requirements include:

- ~180 Watts
- ~8–10 Amps @ 12V
- ~1.5 Amps @ 5 Vsb (2A peak)

Other requirements include the 4-pin (2x2) connector needed to support the dedicated CPU VR input.

For more information, see the *SFX/SFX12V Power Supply Design Guide.*

Summary

The combination of Intel's new generation of Pentium 4 processors in the mPGA478 configuration creates new opportunities for OEMs and system integrators. The combination of new Pentium 4 processors with the microATX form factor means that more processing performance can be made available to consumers in smaller chassis configurations than ever before.

The mPGA478 specification is intended to support Intel's Pentium 4 processor roadmap well beyond 2 GHz. System builders should move their performance PC designs to this configuration in order to take maximum advantage of this stable platform.

The new 478 Thermal Mechanical solution simplifies the assembly process by eliminating direct-attach mounting screws and stand-offs in favor of pushpins. This assembly serves the dual purpose of cooling the processor and chipset while contributing added structural stiffness to the system board. Smooth integration is largely a matter of understanding a set of well-defined requirements for chassis fix, cooling, and the power supply.

While some new components are required, successful integration, as always, is a matter of knowing the requirements, integrating the new technology, and then thoroughly testing the system.

Contact your Intel representative for information on component availability.

More Info

For information on the guidelines and components mentioned in this article, refer to the following sites:

- Information regarding the Intel® Reference Thermal Mechanical Solution is available on the Web. Visit the Desktop Form Factors organization site for the *Performance MicroATX Design Guide 2.0*. Consult your Intel representative for information on where to obtain individual components of the thermal/mechanical solution.
- For more information on meeting Pentium 4 processor power requirements, see the SFX/SFX12V Power Supply Design Guide at the Desktop Form Factors organization site.
- Refer to the guideline entitled Performance microATX Thermal Design Suggestions version 1.0 to understand requirements specific to your configuration. It can be found on the Desktop Form Factors organization site.
- Intel Active Monitor is also available for download from the Desktop Boards area of the Intel Developer Site.

All dates specified are target dates based on current expectations, provided for planning purposes only, and are subject to change.

Author Bio

As a technical marketing engineer in the Intel Architecture Marketing Group, Dan Ragland is responsible for supporting the technical aspects of new Intel desktop board products and technologies. Dan is currently a member of Intel's product development team working on the Intel Desktop Boards D850MD and D850MV. In addition, he supports various emerging desktop board technologies. Dan came to Intel in 1998, joining the motherboard development operation.

Initiatives and Technologies

Intel Press Unveils Four New Technical Books

Richard Bowles
Publisher
Intel Press
Intel Corporation

Overview

At the Fall '01 Intel Developer Forum Intel Press announced four new technical books aimed at streamlining the development and deployment of systems and software applications based on the latest Intel® Architecture products and technologies.

Authored by core technology and product development team members, Intel Press books are among the earliest and most comprehensive means for developers to learn the inner workings of leading computer technologies.

Intel Press has gathered leaders in the field to open their notebooks and transmit their knowledge to the technology community. According to Intel Press, the only way to gain more direct access with industry experts would be in a face-to-face meeting—but even that couldn't match the comprehensive information that these technical books contain.

New Book Titles

The new Intel Press books just released include:

- *Programming Itanium™-based Systems*—The essential programming book for developers of workstation and server applications running on Itanium™ processors
- *InfiniBand* Architecture Deployment and Development*—Lessons on creating, optimizing, and tuning I/O to avoid bottlenecks
- *Building the Power-Efficient PC*—The first comprehensive hardware and software engineering guide for designing power-managed PCs
- *Peer-to-Peer Computing*—A comprehensive explanation of the new alternative to client-server architecture

All four books are detailed resources for those who create or deploy the leading edge of technology. Content ranges from fundamentals like schematics, source code, project files, and step-by-step guides to explanations of the strategy and concept behind the technology.

Programming Itanium™-based Systems

Every software developer targeting Intel's new 64-bit Architecture will gain a world of general programming knowledge in Programming Itanium™-based Systems. Former Intel Application engineer, educator and author of *Itanium™ Architecture for Software Developers* Walter Triebel joins architecture, assembly language, and embedded systems expert Joseph D. Bissell, and performance computing specialist Rick Booth to provide practical software tool and applications information. Included is advice on code optimization, lessons on taking advantage of Intel® Itanium™ Architecture parallelism, and insights into optimizing applications on Intel Itanium processors.

InfiniBand* Architecture Deployment and Development

Developers of software applications and peripheral hardware looking for an edge in system I/O or interprocess communications should have close at hand InfiniBand* Architecture Deployment and Development. The author, Bill Futral, shares insights and experience as co-chair of the InfiniBand Trade Association's Application Working Group and as a seasoned I/O architect on Intel's Advanced Components Division Architecture Team. Also valuable for data center IT executives, the book illustrates ways to implement and deploy InfiniBand to relieve existing I/O bottlenecks.

Building the Power-Efficient PC

This is the first comprehensive hardware and software engineering guide for designing power-managed PCs. Building the Power-Efficient PC is an indispensable guide for all PC hardware and software developers working to meet new ACPI power-management guidelines. Intel system architects and software engineers Jerzy Kolinski, Ram Chary, Andy Henroid, and computer systems designer Barry Press cover the full gamut of power efficiency. Subjects include background and history, specific explanations of hardware and software functions that increase efficiency through low power consumption, automated off-hours maintenance, and improved PC restarting activities.

Peer-to-Peer Computing

Written to establish a general knowledge of Peer-to-Peer (P2P) technology among application developers, IT professionals, and end users, Peer-to-Peer Computing takes readers from early experiments to the current wave of music-oriented applications and into future business and consumer uses. Principal topics are how P2P provides secure direct exchanges, relevant content discovery/delivery, and sharing of resources at the edge of the network. David Barkai, a vector processing pioneer and expert in computer applications and architectures, explains the foundation technologies of P2P and the viability of P2P as an alternative to the client/server architecture for a network-based computing, as well as for self-managed online communities.

Summary

With these four new titles, Intel Press expands its collection of first-to-market resources providing engineer-to-engineer technical information on the most recent Intel developments. Each book offers background on the underlying technology as well as practical directions and instructions for developing and optimizing computer systems for Intel architectures.

More Intel Press publications are under development, including comprehensive books on Intel® Pentium® processor tuning, Itanium processor tuning, scientific programming on Itanium processors, and programming the Intel® IXP1200 network processor family.

Intel Press titles available now include:

- *USB Design by Example, Second Edition* by John Hyde
- *Itanium™ Architecture for Software Developers* by Walter Triebel
- *InfiniBand* Architecture Development and Deployment* by William T. Futral
- *Programming Itanium™-based Systems* by Walter Triebel, Joe Bissell, and Rick Booth
- *Building the Power-Efficient PC* by Jerzy Kolinski, Barry Press, Ram Chary, and Andy Henroid
- *Peer-to-Peer Computing* by David Barkai
- *Building a Simple Network* by Ken Denniston

More Info

Intel Press books can be purchased directly from online retailers Fatbrain, Powell's, and ShopIntel.

To learn more about Intel Press and its publications, visit the Intel Press site. You can also visit the Intel Press Cafe area at Intel Developer Forum events being held around the world this fall.

Author Bio

Richard Bowles is the founding publisher and program director of Intel Press. He has over 15 years of Intel tenure and has served in a wide variety of market, product, and sales development capacities. Since rejoining Intel in 1993, Richard planned and directed the initial phases of Intel's market development efforts in the workstation arena.

Richard has also held senior marketing and general management positions at Hewlett-Packard and Tektronix, and has consulted in market development to several start-up and Fortune 100 companies during more than 20 years in the electronics industry. He holds a Bachelor of Science degree in mechanical engineering from Stanford University, where he also completed post-graduate studies in engineering management.

Creating a Third Generation I/O Bus

Gerald Holzhammer
Director
Technology and Research Labs
Intel Corporation

Overview

Looking back at the original development and evolution of the PCI Local Bus specification provides some interesting insights into our present-day effort to create its successor, the Third Generation I/O Bus (3GIO).

As pervasive as PCI has become over the past decade, it's easy to lose sight of the fact that PCI was initially envisioned as a local chip-to-chip bus only. However, it soon became clear that this initial goal was far too limiting. We came to realize that what the industry really needed back in 1990 was a general-purpose replacement for the ISA bus. As it turned out, that was just the beginning. During the 1990s, PCI evolved from its original role as a general-purpose I/O interconnect for PC clients to become a pervasive I/O backplane in servers, communications platforms, and, as Compact PCI, in a multitude of embedded applications. None of the original development team imagined how widely it would ultimately be adopted.

With 3GIO the industry has the opportunity to learn from the past. By comprehending these requirements from the very beginning, we have an opportunity to create a general-purpose I/O interconnect that provides a unified foundation for multiple market segments. The vision is a single interconnect that can provide the cost-effective performance required in desktop systems, the power efficiency needed for mobile systems, the high performance required in servers, and the versatility to meet the I/O needs of platforms ranging from desktop PCs to Internet infrastructure equipment. Broad adoption will benefit all segments by reducing costs and ensuring reuse.

The goal is to balance the requirements of multiple computing and communications segments while keeping the overall architecture clean and simple. The result is a third generation I/O interconnect that combines the highest available bandwidth per pin with the ability to cost-effectively scale toward the limits of copper signaling technology.

Longevity

Major I/O interconnects, such as ISA and PCI, typically thrive for 10 years and then slowly fade away over time. This longevity has to be comprehended upfront, and it makes scalability a key requirement for any new I/O standard.

The PCI bus has served the industry well since 1990, and it will continue to play a key role for several more years. As an open specification, PCI has been the foundation of continuous innovation throughout the industry. But the inescapable fact is that the evolution of computing is driving I/O bandwidth requirements beyond the capacity of PCI 2.2, PCI-X, or any multi-drop parallel bus architecture.

A key indicator of a technology generation nearing the end of its useful life is the emergence of point solutions, which add platform cost and reduce reuse. In the mid 1990s, interconnect requirements began to exceed the capacity of PCI. As a result, I/O that was once handled by PCI began to migrate to a variety of point solutions. The first, driven by 3D graphics, was the Accelerated Graphics Port (AGP) in 1995. By 1998, increasing demands of CPU and memory I/O led to the creation of Intel's Hub Architecture replacing PCI as the chipset mezzanine bus. Similar examples can be found in the server, communication, and embedded spaces.

Generations of newer and faster processors are emerging at a nonstop pace. The processor system bus continues to scale in both frequency and voltage, and memory bandwidths have increased to keep in step with the evolution of the CPU. Gigabit Ethernet* and InfiniBand* Architecture are driving higher requirements for external I/O bandwidth.

The result is that what once happened to ISA back in 1989 is now happening to PCI. Parallel bus technology is reaching the limits of cost-effective scalability. Attempts to push these limits in order to create a higher bandwidth, general-purpose I/O bus will result in large cost increases for little net performance gain.

Scaling to the Limits of Copper

3GIO is being designed with long-term scalability in mind. Key features are higher bandwidth per pin, low overhead, low latency, and embedded clock architecture. 3GIO is fully serial Interconnect with links implemented using multiple, point-to-point connections called lanes. 3GIO's initial speed of 2.5 Gbyte/sec/direction provides two unidirectional 200-MByte/sec communications channels that represent roughly three times the speed of classic PCI. By adding more lanes, bandwidth can be easily scaled.

Its embedded clock timing and differential signaling enables 3GIO performance to scale to the limits of copper signaling, which we expect to be in the 10–15 Gbyte/sec range for high-volume copper PCB technology. The embedded clock architecture requires fewer pins than parallel architectures, which simplifies routing. This helps minimize costs and allows wider flexibility for component, motherboard, adapter, and system design. Fewer signals also means that systems can be designed using less board space and smaller connectors, which in turn supports smaller and more innovative form factors.

Building on PCI and InfiniBand* Architectures

3GIO has its roots in PCI and InfiniBand architectures. 3GIO adopts the PCI software model. It can boot existing operating systems without changes, and it supports PCI-compatible configuration and device driver interfaces. The built-in stability of 3GIO's software ensures that 3GIO can take advantage of the evolution of physical layer characteristics as new silicon and motherboard technologies evolve without impacting higher levels of the system architecture.

While 3GIO will replace PCI over time, we expect InfiniBand architecture and 3GIO to coexist. InfiniBand architecture's primary focus is on shared I/O in a multi-computer environment using a robust message passing architecture; 3GIO's focus is cost-effective local I/O using a load/store architecture. In addition, 3GIO is building on InfiniBand architecture's electrical layer and silicon building blocks. We see this as a key enabler toward further convergence of the two I/O architectures in the future. One of the scenarios we are exploring is the reuse of a 3GIO device within an InfiniBand fabric.

Industry Working Together

Naturally, to establish something as ambitious as a third generation I/O specification requires broad industry cooperation. We are very pleased that the Third Generation I/O Work Group, consisting of founding companies Compaq, Dell, IBM, Intel, and Microsoft, and the PCI-SIG, a nonprofit special interest group, are working together with many other key companies to define this new serial I/O interconnect. The Work Group has completed a preliminary specification and is on-track to deliver a specification to the PCI-SIG for review in the first quarter of 2002.

Summary

CPU speeds in excess of 10 GHz, faster memory speeds, higher-speed graphics, 1 Gigabit and 10 Gigabit LAN, 1394b, and other technology advances will drive I/O requirements beyond the current capabilities and cost-effective scalability potential of the PCI bus. 3GIO is extremely important to developers because it creates a high-performance, highly scalable, general-purpose I/O architecture.

This new architecture is being developed jointly by the Third Generation I/O Work Group and will be made available for broad industry adoption through the PCI-SIG. Additionally, it will be designed to serve as a long-term, general-purpose I/O interconnect to meet the requirements of desktop, mobile, server, communications, embedded, and other emerging and future applications.

Developers should plan investments now in order to have products based on 3GIO architecture ready by the second half of 2003.

More Info

For more details on 3GIO architecture, read the white paper *Creating a Third Generation I/O Interconnect* by Ajay Bhatt (PDF, 1.15MB).

Author Bio

Gerald Holzhammer manages Intel's Desktop Architecture Lab, the group tasked with research and development of next-generation desktop PC architectures, interconnects, and platform ingredients including power delivery, thermal, EMI, and acoustic solutions. His group has also been involved in the enhancement of PC ease-of-use and the development of advanced proof-of-concept platforms.

Since joining Intel in 1990, Gerald has managed several architecture, multimedia, and desktop labs with a focus on PC platform innovation. He led the architecture teams that developed the PCI, USB, and AGP specifications, which became widely adopted by the PC industry.

Before joining Intel, Gerald worked at Siemens AG Germany/USA, where he worked in a joint venture between Intel and Siemens to develop 32-bit object-oriented platforms. He also worked at BiiN, USA, where he managed the kernel team for an advanced object-oriented multiprocessor system.

Gerald received his M.S.E.E. from the Technical University of Vienna, Austria, and he holds a B.S.M.E. from HTL Innsbruck, Austria.

An Integrated Approach to Audio Design

Fred Loeb
Applications Manager
Computer Products and Services Group
Analog Devices, Inc.

Overview

Audio system designers face several tradeoffs in creating solutions for high-quality, low-cost subsystems for PCs and consumer electronics devices. None of these tradeoffs is completely satisfactory in achieving exceptionally high-quality audio reproduction at relatively low costs, a goal that is increasingly important in today's brutally competitive electronics markets.

These tradeoffs become more difficult in the current design environment, where PC manufacturers attempt to improve the quality of user experience while reducing engineering and component costs. In most cases, this requires reducing the size of the BOM (bill of materials) altogether. Rather than simply designing to a given level of audio quality, engineers are being asked to build in both high audio quality and low cost.

Rather than submit to these design tradeoffs, an Analog Devices solution combines aspects of a traditional AC'97 codec with features of a DSP to produce an integrated audio subsystem that can provide lower overall costs while improving audio quality. This solution is also software-configurable, which provides engineers with greater flexibility in tuning audio subsystems for specific speakers or audio conditions.

Today's Design Choices

Three different design alternatives are representative of the types of tradeoffs engineers face in creating audio subsystems. Think of these three alternatives as points on a continuum of choices for OEM designers, using cost of implementation as a measure.

1. Audio reproduction with no on-board equalization, compensating for reduced audio quality with the requirement for expensive speaker/headphone technology. The lack of parametric speaker equalization on the audio reproduction hardware results in audio reproduction that is inflexible and will probably not produce the desired audio qualities, especially when used with inexpensive speakers. The resulting output cannot be adapted by engineers for a particular set of speakers or headphones, and cannot be adjusted by end users. Because the output can't be adapted to the audio reproduction hardware or the acoustics of the surrounding environment, it is highly unlikely to be identified by the human ear as a true reproduction. Some frequencies will appear to be highlighted, while others will sound subdued. The audio will be reproduced, but not in its seemingly original form.
2. The use of a standard AC'97 codec such as the Analog Devices AD1886 with the addition of a discrete analog equalizer implementation, which supports higher quality audio reproduction using less expensive speakers. The use of parametric speaker equalization as a part of the audio subsystem means the same sound at a lower cost, or better sound at equivalent cost. While the use of analog linear equalization on the audio subsystem improves quality, it comes at a cost. To perform analog equalization requires a significant level of complexity on the subsystem. In particular, each bandpass filter requires an op amp, two capacitors, three resistors, and other support components. Analog equalizers must also be tuned manually, which requires expensive experience in that practice. With less experienced engineering resources, the result could still be lower quality audio.
3. A high cost solution combining two chips such as a DSP/codec chip integrated with an AC'97 codec. This solution incorporates full digital parametric equalization with the incorporation on a DSP into the design. The digital DSP design alternative requires an extra codec to convert the analog signal to digital for processing, plus a separate DSP, in addition to a number of support components. This represents an example of the highest quality available today in an audio design, but at a high cost.

Any of these types of designs work well enough for today, and are used extensively, but they're no longer fully satisfactory. The dynamics of the PC and consumer electronics industries continue to demand both high quality and low costs. Yet design expertise to fully achieve this goal is often lacking. Audio engineers require better alternatives for lowering overall system costs while improving sound quality.

The Advantages of an Integrated Approach

The traditional design curve for audio reproduction can be redrawn by offering a higher level of electronics integration onto the codec itself. This approach can address the issue of the use of lower quality support components in the manufacturing process by simply eliminating the need for an extended design incorporating op-amps, resistors, and capacitors, while achieving the same effects on-chip.

This approach also addresses the variability that can be introduced at manufacturing time, while offering the potential to control costs by replacing discrete components with an integrated chip. This, of course, assumes that the integrated chip can be manufactured and made available for less than the total cost of the several support components, which is usually a reasonable assumption.

The integrated approach can also take aim at improving the quality of audio reproduction at the design stage. This is the design goal of the Analog Devices AD1981.

AD1981 Design

The AD1981 is a combination of a fixed AC'97-compatible codec with DSP features. It offers significant design advantages over any of the alternatives available. Today, the DSP features enable adjusting the center frequencies and bandwidth, to be able to tune for different speaker sets. And because the DSP features are integrated with the codec, the manufacturing cost of the subsystem can be reduced without compromising on audio quality.

The AD1981 provides OEM audio engineers with the ability to reduce costs and improve audio quality and manufacturability while providing them with the opportunity to increase the range of design options available. It offers significant improvement in cost by requiring fewer components to support a mid-to-high level of audio quality, while enabling engineers to achieve better control of acoustic design and audio quality. From the standpoint of manufacturing, it provides the benefits consistent with the expectations of a separate DSP solution—that is, a design free from component tolerances and drift—while providing the ability to do circuit tuning to improve audio output.

The AD1981 has added features in addition to the AC'97 native device drivers that enable OEMs to deliver the added value of an AD1981-based audio subsystem. Because enhanced drivers are provided with the AD1981 itself, OEMs will find that the lowest cost solution will be to provide their customers with these drivers rather than writing their own. However, if the AD1981 driver is unloaded or unavailable, the subsystem will still produce quality AC'97-compliant audio output.

There is also a method of adapting filtering in software during the audio subsystem design. The software used for this process incorporates a graphical user interface that presents the engineer with a visual view of the filters at a given set of frequencies. With this tool, engineers use a text box to assign the center frequencies and bandwidth from which to set the filter. The attenuation at that frequency band can then be adjusted through a graphical slider control. Once the adjustment is made, the engineer receives feedback in the form of a graphical image of the resulting frequency response for a given output.

This affords the audio subsystem designer the opportunity to tune the output of the AD1981 specifically for a given set of speakers. For example, the design engineer can use the AD1981 software to build a frequency response curve that exactly complements the speakers, producing a flat response across the audio spectrum. This process may take no longer than a few minutes, yet it can result in audio reproduction significantly better than if there were no knowledge of the speakers used.

Summary

Using AD1981's DSP-like capability to provide parametric speaker equalization means the same quality audio at a lower cost, or better audio at equivalent cost to existing equalization solutions. Audio subsystem designers are able to use the device to break the traditional design parameters to deliver quality audio for PCs and other consumer electronic devices at prices consistent with OEM and end-user expectations.

The AD1981 is the result of a continuing commitment by Analog Devices to serve the PC market that monitor and manage power usage, process signals used in flat-panel displays and multimedia projectors, and enable PCs to provide CD-quality audio. Analog Devices also serves the high-end consumer market with integrated circuits used in products such as digital cameras and camcorders, DVD players, and surround-sound audio systems. In the industrial market, Analog Devices provides a broad portfolio of components, including products for automatic test equipment and for the digital speed control of AC motors.

More Info

More information about the AD1981 can be found on Analog Devices' Web site, or contact soundmax@analog.com.

Author Bio

Fred Loeb is applications manager in the Computer Products & Services Group at Analog Devices, Inc. His background combines a focus in new product definition for analog ICs working with PC OEMs to deliver new Computer audio products on both motherboards and sound cards. He was awarded a B.S.E. from SUNY, Stonybrook. He has furthered his technical knowledge with advanced education in electrical engineering from Tufts University (Medford, MA).

New Technology Brings RosettaNet* to Smaller Businesses

Shuvranshu Pokhariyal
Senior Software Engineer
Intel Architecture Group
Intel Corporation

Overview

RosettaNet* is an e-Business standard that facilitates automated transactions between businesses. It standardizes the XML payload of the message and the process that goes with it.

RosettaNet is a complex standard that until now has been expensive to implement. In fact, current RosettaNet solutions can cost millions of dollars, hindering the proliferation of e-Business.

Intel is committed to accelerating the adoption of the RosettaNet standard. Its e-Business Group is a champion in promoting RosettaNet standards inside and outside Intel, and the company sits on the board of the RosettaNet Consortium. Now Intel is backing up its commitment by creating supporting technology focused primarily toward small and medium enterprises, where Intel market analysis shows adoption of the RosettaNet standard is slow.

Intel® Whitehorse technology, created by the Intel e-Business Group, is a cost-effective and efficient solution that enables RosettaNet transactions expeditiously for small, medium, and large enterprises. Whitehorse technology enables automated business transactions in a very affordable way, without proprietary networks or satellite products. It is an open-standard-based implementation that works with any technology following RosettaNet standards. While aimed at small-to-medium businesses, it scales well to the needs of larger companies.

Purchase Order Management

The RosettaNet Standard defines Partner Interface Processes (PIPs) that it groups into eight clusters: RosettaNet support, partner product and service review, product information, order management, inventory management, marketing information management, service and support, and manufacturing. Whitehorse currently implements the six PIPs in the purchase order management cluster (manage purchase order, notify of purchase order acceptance, notify of advance shipment, notify of invoice, notify of invoice reject, and notification of failure). Intel has determined that these six are the most useful and important for getting businesses up and running with RosettaNet.

Whitehorse technology can be used as a developer toolkit or as an end-user product. The end-user component has user interfaces that make it quite simple for a business to automate its purchase order management. For instance, a small business can construct a purchase order by simply filling out data on a screen and hitting the submit button, all according to the RosettaNet standard, so that another business can automatically process that request.

For larger companies that require back-end information and need to extract data from a large database, Whitehorse includes user interfaces that provide the ability to specify locations for extracting data, enabling back-end integration. For instance, if a business wants to construct a purchase order for products from a large company, all the user needs to do is specify line items. That particular data would be loaded automatically, and the user simply hits the submit button.

Whitehorse simplifies the management of trading partners with a simple and easy-to-use user interface. It also supports the concept of major and minor trading partners, allowing major trading partners to specify the usage of optional fields, and other process customizations. (A major trading partner might typically deal with hundreds or thousands of businesses, while a minor trading partner may be a small mom-and-pop shop.) The customizations of PIPs can then be accepted and installed by the minor trading partner.

Another important feature for end users is that Whitehorse makes it easy to support Digital Certificates, which are needed to establish a connection via SSL (secure sockets layer), and for digital signatures and data encryption.

Designed to Meet Current and Future Needs

The developer toolkit side of Whitehorse technology includes three components: user interface and database integration, RosettaNet standard implementation, and communication protocols.

The back-end and database interface provides an API that developers can use when they have workflow-type applications, or if they already have their private processes well established. For instance, a purchase order could be automatically triggered to a particular person via the RosettaNet standard. Since the product is written in Java*, the APIs are currently RMI-based (Remote Method Invocation-based), allowing distributed objects (on different machines) to use the interface. In future releases, SOAP-based (simple object access protocol-based) interfaces will be provided.

Developers who simply want to use the "RosettaNet standard component" alone will find both RNIF 1.1 and RNIF 2.0 of the RosettaNet Interface Framework are supported. This component handles all the standard's requirements, including process and state management, automatic receipt acknowledgements, non-repudiation of data, and encryption.

The communication protocols component abstracts the details of the underlying implementation providing an easy-to-use interface that includes the protocol to use, the URI (universal resource identifier), and the data to be transported. Supported communication protocols currently include HTTP and HTTPS/SSL as mandated by the RosettaNet standard. There will soon also be support for SOAP and SMTP (simple mail transfer protocol).

While Whitehorse currently supports only six PIPs, it is designed so that it can accommodate new PIPs very easily. An automatic mechanism currently in development will make it possible for a developer to add PIPs at will.

While the Whitehorse toolkit is currently Java-based, future releases will be language-agnostic with SOAP support. Future versions of Whitehorse will support the ebXML* Messaging Services Specification and UDDI. Note that ebXML is a set of specifications that, together, enable a modular electronic business framework. UDDI, the Universal Description and Discovery and Integration project, has a global repository that creates a platform-independent, open framework for describing services, discovering businesses, and integrating business services using the Internet.

Cost Reduction through Automation

The cost of implementation of the RosettaNet standard is usually a barrier to entry because it could take millions of dollars in development. The standard is quite broad and difficult to implement, with many nuances. In contrast, the cost of implementing Whitehorse is expected to be under ten thousand dollars.

Currently, businesses typically handle transactions by having people fax purchase requests for products to other businesses. This methodology is error prone, introduces inefficiencies in the form of double data entry, and entails unnecessary manual intervention. The estimated cost of these inefficiencies is thought to run in to billions of dollars in the computer industry.

Given the highly competitive nature of the industry, small differences in the supply chain management determine industry leaders. Problems like inventory management have plagued the industry, but automated and standardized e-Business solutions such as RosettaNet promise to change this.

For companies that were already using automated e-Business systems over proprietary networks (such as EDI), the advantage of using the Internet and standardized protocols results in reduced day-to-day operation costs and a much larger group of trading partners. Whitehorse reduces the costs of running the entire system because it makes use of the Internet and an open standard that is XML-based rather than the proprietary networks required by current electronic-based systems.

Summary

RosettaNet is an e-Business standard that facilitates automated transactions that until now have been expensive to implement. Intel Whitehorse technology is a cost-effective and efficient solution that enables RosettaNet transactions expeditiously for small, medium, and large enterprises.

More Info

For more information about the RosettaNet standard, visit the RosettaNet organization site.

For information about Whitehorse contact the author of this article, Shuvranshu Pokhariyal.

Author Bio

Shuvranshu Pokhariyal is one of the original architects of the Whitehorse product. He is a senior software engineer in Intel Architecture Group within Intel Labs. Previously, he worked in the OEM Platform Solutions Division (OPSD). Pokhariyal received his master's degree at the Oregon Graduate Institute.

Creating Peer-to-Peer Middleware from Web Services Technologies

Cedric Yau
Product Manager
Solutions Enabling Group
Intel Corporation

Overview

Over the past year, Peer-to-Peer (P2P) and Web Services have become popular as viable computing models. Both are examples of a Service Oriented Architecture (SOA) and have many similarities in underlying technology. They also have several differences arising from the intermittent connections and dynamic network locations found in Peer-to-Peer systems.

Recent advances in open standards and technology frameworks have made development and deployment of Web Services faster and easier. Because of the similarities between Web Services and Peer-to-Peer communication, Web Services frameworks can also be used as a basis for Peer-to-Peer development. However, there are important differences in the Peer-to-Peer programming environment that create the need for additional enhancements.

This article will discuss the similarities and differences between Peer-to-Peer and Web Services and introduce Intel's Peer-to-Peer technology building blocks initiative.

Peer-to-Peer Background

At its core, Peer-to-Peer (P2P) computing is the sharing of computer resources and services through direct communication between systems. Computers that traditionally acted as clients now incorporate server capabilities that enable them to share processing power, bandwidth, and storage.

This form of computing can reduce the network load utilized by central servers, enabling larger groups of users to be served by the same infrastructure. There are several use models for P2P computing. For example, P2P can be used to enable rich collaboration. While instant messaging works well in a client/server environment for plain text, it breaks down with higher bandwidth forms of communication such as voice and video. Real-time collaboration products such as Microsoft NetMeeting* and Groove Networks Groove* application can shorten the time it takes to complete tasks by enabling people to communicate remotely.

Collaboration may also occur between intelligent agents. These agents communicate with each other to benefit their end users. Previous work done by Sandia National Labs has demonstrated the ability of P2P agents to detect and prevent virus threats from spreading across a network. Intelligent agents can also work together in the work environment to facilitate data sharing. Enterprise-wide searches can be performed to locate knowledge usually hidden in documents stored on personal computers.

Finally, P2P programming may be used to distribute tasks around a network. This can lower the cost-of-entry for businesses to utilize large-scale computing needs such as simulation, rendering, or statistical analysis by distributing tasks to idle computers around the network. Rather than having to purchase or lease time on a dedicated cluster, companies can utilize existing employee desktops to perform computationally expensive tasks.

Web Services Background

Web Services comprise remotely accessible software components that are utilized by multiple parties. They are an extension of object-oriented design that enables cross-organizational solution development through the creation of interoperable building blocks.

Web Services serve as the basis for higher-level capabilities and typically perform single tasks such as order processing or returning search results. More importantly, Web Services utilize existing Internet protocols and can be accessed without tight integration as with previous RPC (Remote Procedure Call) architectures. In addition, Web Services are able to describe their own input and output requirements, allowing quick and dynamic configuration and re-configuration.

Web Services' standards, most notably XML, HTTP, and SOAP, are of interest to Peer-to-Peer developers. XML provides a universal data format, HTTP, a standard communication protocol, and SOAP, a standard interface for object RPC. These standards have led to the development of frameworks and tools for creating and deploying Web Services.

Creating Common Ground

In the past, P2P developers had to create their own application stacks from the ground up. For most, this required the definition of communication protocols, implementation of a communications stack to resolve messages to function calls, and creation of multi-threaded server code to handle simultaneous requests.

This form of development is generations away from what exists in the Web Services market today. Tools and frameworks today enable developers to efficiently enable class object for remote access through SOAP/HTTP, with little more than the addition of a few tags and the running of a few commands. Additional tools are also offered to enable remote objects to seamlessly be integrated into local applications.

Fortunately, the standards and frameworks used to create Web Services can also be utilized to develop P2P applications. This is because both sets of architectures fundamentally coordinate interactions between loosely coupled systems. Utilizing a common framework based on current Web Services technologies would enable P2P developers with elementary building blocks for building applications. This would reduce development time and eliminate the need to learn multiple development platforms.

Differences Between P2P and Web Services

Web Services frameworks, however, are not a panacea for Peer-to-Peer development. On the surface, one major difference is that P2P applications are stand-alone applications and not hosted from a Web server. More importantly, P2P offers new challenges not found in the Web Services environment. For example, peers may drop off at any time. There is no guarantee of 99.999 percent uptime or of any availability, for that matter. At a higher level, this means peer presence must be tracked since it cannot be assumed as with Web Services. At a lower level, the lack of availability also creates the need for a reliable messaging facility with an option for redelivery at a later time. Also, unlike Web Services that are addressable at well-known hosts and ports, peers may alter their location on their network. This creates the need for an ability to resolve peer names to network addresses. Finally, while Web Services are hosted on external servers and addressable from outside the network, peers are often hidden behind NATs (network address translations) or firewalls and may require the use of a tunneling protocol or relaying services for communication.

Introduction to Intel's P2P Technology

Peer-to-Peer middleware can be created within a Web Services framework. With its highly integrated set of tools for Web Services and application development, the Microsoft .NET* Framework provides an excellent platform to begin. In order to enable robust Peer-to-Peer capability development, building blocks supporting location independence, security, and availability are needed. As a result, Intel is creating technology to facilitate development of P2P applications based on the .NET Framework.

One aspect of Intel's initial technology is based on the Remoting feature of the Microsoft .NET Framework. Remoting is a powerful and easy to use mechanism for interaction between .NET application domains and is a key element of Web Service support in .NET. It automatically formats network messages and remote calls using the SOAP standard, making it an attractive interoperable basis for P2P communications.

Remoting abstracts the complexities of calling remote methods. When a remote object is called, clients create proxy objects that serve as representatives that forward calls to the remote server. Messages are transported to and from remote objects through a channel.

Remoting works well for exposing objects as Web Services. Unlike Web Services, however, peer systems exist in a more transient and difficult environment. Often they do not have a unique or known DNS name. They are not always on, not always connected to a network, and when connected might not always have the same IP address. Often on their network they are behind a NAT box or a firewall, so that outgoing connections work but incoming connections, as from other peers, are blocked. Thus, while remoting works well for Web Services, it is not well suited for P2P applications in constricted environments.

Intel's Peer-to-Peer technology builds upon Microsoft's .NET Framework in several ways. First, it adds a new peer scheme URL to identify objects to remoting that offers a namespace that gives all peers a consistent, durable, and unique name. Essentially, it is a GUID (Globally Unique Identifier) wrapped up to be syntactically DNS-compatible. It also allows parameters to configure how peer communications should work per object. This namespace is integrated into remoting and allows objects to be referenced through a peer URL.

Another aspect of the technology is a single listener that can be shared across multiple applications. All incoming P2P remoting calls pass through the listener, which forwards them to the appropriate application. The target application is denoted in the URL. At initialization time the listener establishes a set of regular HTTP and HTTPS scheme URLs that can be used to connect to the listener.

Security and Availability Enhancements

SSL will be added for simplistic authentication and encryption on the network without object hosting through Microsoft IIS. The authentication approach allows peer computers to recognize each other via a single certificate established by a user of the peer. Between such peers any application can interact with any other application using SSL encrypted communications.

If an application has information to send to another peer and doesn't need to do so interactively, it can enable the optional store-and-forward service when it creates a proxy object. This works with one-way remoting calls. If the destination peer is not online when a remoting call is made, the call will be stored and periodically retried until the destination peer comes online or the call times out.

Finally, if the peer is behind an impenetrable NAT or firewall, it can be configured to use a relay server. The listener calls out to the relay server and establishes a BEEP (Blocks Extensible Exchange Protocol) tunneling connection. The relay server can then accept incoming connections from the outside and pass them through to the listener on BEEP channels. The listener establishes with the relay server the external URL to be used and adds it to the list.

Once the listener has established a list of URLs, it registers them with a directory service under its peer name. In the future, other NAT, firewall, and related solutions can be implemented that cover more situations and might generate additional listener URLs.

With Intel enhancements a developer can use .NET remoting in P2P applications without concern for some of the peer environment difficulties of intermittent connectivity, naming, NATs, and firewalls. The enhancement package can be extended in the future with more solutions to network problems and with richer semantics for security and availability operations.

Simplified File Sharing

Many existing P2P applications are built around sharing or distributing files. Collaborative applications often include mechanisms to share files while content distribution applications utilize the storage and network connectivity of peers more efficiently than centralized servers.

In addition to adding P2P technology to Microsoft's .NET remoting, Intel's technology will add an easy-to-use facility for copying files, based on the remoting enhancements. It is designed to allow users to maintain control over their files while allowing developers to quickly add file sharing and transferring capabilities. Files are transferred in chunks. Since a large file may take a large amount of time, transferring a file in chunks with a simple retry procedure enables the transfer to continue in the event a chunk fails.

The File Copy API is designed to ensure that an application at the source has access rights to read a file, and an application at the destination has access rights to create and write the file. This prevents an application or user from copying files without proper permissions.

Another aspect of the API design involves the form of the URL, which acts as a handle to the file. The URL does not reveal the original location of the file, and instead contains a random value that is difficult to guess. This enables the application developer at the source end to implement control, on behalf of the user, over what other peers get the URL and are then able to copy the file.

Configurable source and destination repositories give the user a second point of general control over files. A user specifies repositories independently of applications so sensitive directories can be protected from accidental or intentional misuse by any P2P application.

Summary

Over the last year, Peer-to-Peer computing has grown in significance. Use models have been established for collaboration, content distribution, and intelligent agents. Historically, such applications have developed from scratch, using proprietary protocols and implementations. Work was also duplicated between solutions due to the lack of available building block technologies.

Recently, the Web Services space has created tools and frameworks based on industry standards such as XML, HTTP, and SOAP. Due to the similarities between Web Services and P2P, Web Services tools and frameworks can also be utilized as a basis for P2P. However, there are important differences between the two, arising from particular P2P needs, such as dynamic addressing and intermittent connectivity.

Intel is creating building block middleware based on Web Services technology to address these differences. Release availability will be announced in the fourth quarter of 2001. This will enable developers to incorporate new P2P capabilities into existing solutions and create new stand-alone applications.

More Info

For more information, visit the Peer-to-Peer Initiative section of the Intel Developer Services Web site.

Author Bio

Cedric Yau is a product manager for the company's P2P technology initiatives. Cedric has also worked on the eBusiness Architecture Marketing team, evaluating potential strategic technologies. Prior to joining Intel, Cedric received a B.S. and master's in computer science from the University of Illinois at Urbana-Champaign and participated in a P2P content distribution startup.

Networking & Communications

Intel® WAN/LAN Access Switch Example Design Highlights the IXA Ecosystem

Bryan A. Brooks
Technical Marketing Engineer
Network Processor Division
Intel Corporation

Overview

Developers creating products for communications service providers face ever-increasing demands for longer product life, flexible interface options, and top wire-speed performance. To meet such demands, developers have relied on the Intel® Internet Exchange Architecture (IXA) ecosystem of silicon and software products such as the Intel® IXP1200 Network Processor. Now developers have another tool they can use to gain insight on designing and deploying networks of the future: the Intel® WAN/LAN Access Switch Example Design for the Intel IXP1200 Network Processor.

The Example Design provides a framework for custom designs by demonstrating an entry-level access switch using Intel® IXA building blocks to develop next-generation networking applications. Using products that are available from Intel and Intel IXA Developer Forum members, the Example Design illustrates the use of several components including the Intel IXP1200 Network Processor, Linux* with the OS-agnostic ATM (Asynchronous Transfer Mode) subsystem software, and the Microware* Microcode Solutions Library* AAL5 Microcode. Other features of the Example Design include OSPF and RIP routing protocols on StrongARM*, Layer 3 Forwarding and ATM OC-3 to Quad 10/100 Ethernet functionality.

Benefits of Intel® IXA Building Blocks

By providing an application for system-level profiling in a comprehensive, integrated package, the WAN/LAN Access Switch Example Design shows developers how they can expedite product development and reduce component cost by using products and technologies available today from Intel IXA and the open-source community. For instance, the Example Design features software from Trillium Digital Systems and Microware Corporation as well as components from the Linux open source community. Intel enables network developers to get their product to market via several of these options:

4. Buy off-the-shelf silicon and software products available today from Intel and/or Intel IXA members.
5. Use the consulting services of Intel and Intel IXA members to implement custom hardware and software solutions.
6. Allow the customer to implement the design using Intel IXA building blocks such as the IXDP1200 Advanced Development Platform and IXA SDK V2.0.

As shown in Figure 1, the WAN/LAN Access Switch Example Design consists of two major subsystems: the microcode software subsystem (data plane) and the Linux software subsystem (control plane). The microcode subsystem consists of AAL5 microcode for microengines, and feeds packets or cells to the TCP/IP and ATM stacks. The Linux subsystem resides on the Intel® StrongARM* core and hosts the drivers, operating system, and ATM and TCP/IP stacks.

Software Architecture

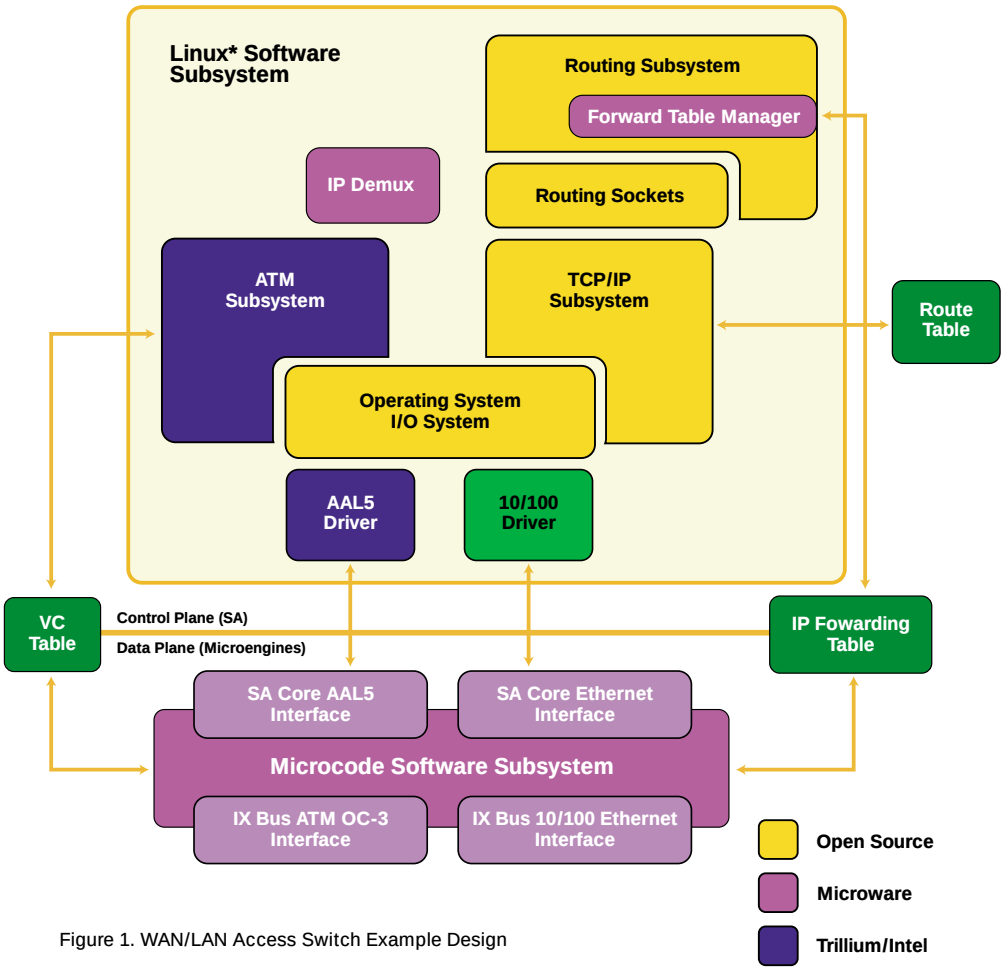

Figure 1. WAN/LAN Access Switch Example Design

Looking Inside the Example Design

To look more deeply inside the Intel WAN/LAN Access Switch Example Design for the Intel IXP1200 Network Processor, consider its primary components:

- *Intel® IXDP1200 Development Platform.* This highly programmable development platform enables designers to define flexible feature sets in software throughout the development cycle, and its easy field reprogrammability supports advanced development. Key components of the development platform that are used in the Example Design are (1) the Intel® IXF6012 Multispeed OC1-12 SONET/SDH Cell/Packet framers, which enable developers to build scalable equipment that can be easily upgraded as the network grows; and (2) the Intel® IXF440 Dual-Speed Multiport Ethernet MAC, which provides a combination of high-performance packet transfer features, comprehensive packet handling capabilities, and network management support.
- *Microware Systems Corporation IXP1200 Microcode Solutions Library v3.0.* This library provides an open, standard interface for fast and easy integration of operating systems running on the Intel StrongARM core, thus enabling developers to focus on product differentiation.
- *Trillium Digital Systems Protocol Software.* The WAN/LAN Access Switch Example Design uses the Trillium Linux SSI, Q.93B, QSAAL, and PLOA software modules. These components enable developers to build a complete suite of communication devices to efficiently pass multimedia traffic over an Asynchronous Transfer Mode (ATM) transport.
- *Intel® Big-Endian Kernel Release.* This is Intel's first release of a Linux port to the IXP1200 and the first release of a Big-Endian Linux port to the Intel StrongARM core. It includes a kernel driver for IXP1200-specific peripherals as well as a RAM disk image for embedded systems development.
- *Intel® Internet Exchange Architecture Software Developers Kit (SDK) v2.0 for the IXP1200 Network Processor.* The SDK provides developers the robust tools needed to fully harness the power of the IXP1200 family and introduces a high-level programming framework that simplifies application development while

reducing time-to-market. The SDK also supports distribution of Linux, which enables developers to reduce costs while benefiting from the extensive Linux ecosystem.

For further detail on the components of the WAN/LAN Access Switch Example Design and its fundamental architecture, see Figure 1.

Summary

The Intel WAN/LAN Access Switch Example Design for the Intel IXP1200 Network Processor clearly shows that Intel has a firm footing in the network processor business. For instance, the Example Design demonstrates that building blocks for IXP1200 designs, such as silicon solutions, software, and hardware components, are available off-the-shelf from third-party vendors, from the open-source community, and from Intel itself. By leveraging the framework offered by the WAN/LAN Access Switch Example Design, developers can take advantage of these building blocks in order to shorten their product's time-to-market and lengthen its time-in-market.

More Info

For more information on the Intel WAN/LAN Access Switch Example Design for the Intel IXP1200 Network Processor, visit any of the following Intel sites:

- Intel Developer Site home page
- Intel Internet Exchange Architecture home page
- Networking and Communications Building Blocks home page
- Intel Literature Center

Trillium refers to Trillium Digital Systems, Inc., the Networking Software Division of Intel.

Author Bio

Bryan Brooks recently joined Intel as a technical marketing engineer for Intel Communications Group, Network Processor Division. Bryan has led the development of marketing collateral, technical presentations, demonstrations, and support functions of the WAN/LAN Access Switch Example Design.

Before joining Intel, Bryan worked at Motorola Incorporated where he worked in wireless technologies as a strategic planner and product marketing engineer for the deployment of 3G wireless technologies such as W-CDMA, CDMA2000, and GPRS. Prior to that, he was a software engineer for six years developing Windows*-based graphical user interfaces for paging and cellular band test equipment and telephony.

Bryan earned his M.B.A. from Arizona State University and holds a B.S.E.E. from University of Maryland, College Park.

Servers

Breakthrough Performance for Intelligent Internet Storage

Wendy Vittori
General Manager
I/O and Bridges Division
Intel Corporation

Overview

Terabytes of data flow across the Internet each day, moving between millions of networked storage subsystems. The surging demand for highly available high-performance storage has shattered the traditional data storage paradigm of local disk drives attached to servers. One of the most profound impacts of the Internet is the migration of storage technologies and subsystems from local servers into the network infrastructure.

The latest generation of Internet storage solutions is based on distributed storage architectures that implement new networking and storage technologies, including InfiniBand* architecture and the emerging iSCSI specification for storage over Internet Protocol (IP) networks.

Intel is involved with emerging storage technologies on many fronts, providing building blocks and tools that enable OEMs and developers to shrink time-to-market while delivering innovation and enhanced product value. From the plug-and-play convenience and scalability of network attached storage (NAS) to high-performance innovations in storage area networks (SAN), Intel is helping the industry meet the demand for intelligent Internet storage with breakthrough performance and high availability.

Intel® storage processors, intelligent storage adapters, RAID-on-motherboard solutions, iSCSI building blocks, storage reference designs, and development tools are helping developers translate their vision of intelligent Internet storage into real-world solutions.

Affordable RAID

The data protection and availability benefits of redundant array of independent disks (RAID) storage has historically been priced beyond the reach of entry-level and midrange server users. The Internet-driven demand for mission-critical data availability makes RAID data protection an increasingly important standard feature in servers at every price point.

Based on its core competencies in silicon, middleware, and integrated RAID solutions, Intel provides a varied choice of flexible RAID building blocks designed for platforms including server-attached storage and external storage arrays, SAN systems, and easy-to-deploy NAS appliances.

Intel storage processors based on Intel® XScale™ microarchitecture provide the intelligence that server OEMs and independent hardware vendors need to support tomorrow's storage performance requirements, while differentiating their RAID solutions. For Internet storage applications, "intelligence" means providing programmable silicon with the power to support the functional requirements of a variety of application environments.

This in turn enables vendors to add value-added software in storage products designed for a variety of competitive price points. The breadth of Intel's family of intelligent Internet storage building blocks enables vendors to build products with multiple levels of performance and cost, tailored for the requirements of these specific market segments.

Building Blocks

Intel® building blocks are at the heart of advanced designs for intelligent Internet storage. They include:

- The industry's most advanced storage processor roadmap, including the Intel® IOP310 I/O processor chipset, the first Intel® product based on the breakthrough performance of Intel XScale microarchitecture.
- Intel® Integrated RAID reference designs and advanced RAID management software.
- Intel® Architecture server platforms and pre-validated RAID-on-motherboard solutions.
- Embedded Intel® Architecture processors, chipsets, and Communications Reference Designs that accelerate time-to-market development of innovative SAN and NAS solutions for intelligent Internet storage.
- An iSCSI host bus adapter that enables data sharing over Gigabit Ethernet networks.
- Comprehensive development tools and software from Intel and a growing list of leading third-party hardware and software vendors.
- A varied selection of best-in-class PCI-PCI bridge devices and Ethernet controllers.

Architectures

Intel provides a choice of processor architectures for networking applications, including Internet storage. Each is optimized to meet specific application requirements:

- Intel XScale microarchitecture made its debut in the Intel IOP310 I/O processor chipset, a high-performance storage processor solution optimized to quickly move large quantity blocks of data in I/O subsystems. The processor features a 733-MHz core and the companion chip provides peripherals that enhance data flow in Intelligent RAID and Internet storage applications.
- The Intel® IXP1200 Network Processor family is specifically designed for network control tasks, such as wire-speed switching and routing of packets or ATM cells in real time. Other control plane functions include exception handling, initialization, and configuration in addition to logging and set-up functions.
- Embedded Intel Architecture processors, chipsets, and other building blocks provide the performance, scalability, and upgradability to meet the application services requirements of next-generation networks. Intel Architecture processors are specifically designed to handle compute-intensive functions required by the application services layer, such as routing and signaling protocols and the control of policy, Quality of Service (QoS), and security.

Breakthrough Performance

The Intel IOP310 I/O processor chipset, including the Intel® 80200 processor and Intel® 80312 companion chip, delivers a combination of breakthrough performance and intelligent integrated features for Internet storage.

- The Intel 80200 processor with Intel XScale microarchitecture provides industry-leading MIPS per milliwatt performance. At 733 MHz, the processor is rated at up to 915 Dhrystone 2.1 MIPs at less than 1.3 watts. The 32-bit processor core includes a number of enhancements, including Intel® Superpipelined RISC Technology, a 7-stage integer pipeline with Dynamic Branch Prediction implemented through a 128-entry branch target buffer, and an 8-stage memory pipeline. Performance enhancements also include a 32-byte instruction cache, 32-Kbyte data cache, and a 2-Kbyte mini-data cache.
- The Intel 80312 I/O companion chip is optimized for the high data throughput requirements of Internet storage applications. The peripheral set includes an integrated 66-MHz 64-bit PCI-to-PCI 2.2 bridge with primary and secondary PCI buses that doubles I/O bandwidth compared to previous-generation I/O processors. An 800-Mbyte/sec PCI-like internal bus connects the core interface unit, 100-MHz SDRAM interface and the expanded 1-Kbyte XOR Application Accelerator Unit used for RAID parity calculations. The companion chip supports up to 512 Mbytes of error-correcting code (ECC) memory.

Choice of Tools

With the introduction of the IOP310 I/O processor chipset with Intel XScale microarchitecture, Intel provides a direct transition path for developers who have been using Intel i960® I/O processors or the first generation of Intel® StrongARM* processors. The Intel 80200 processor supports the ARM instruction set v.5TE, and the Intel 80312 I/O companion chip shares the same I/O mapping and peripheral compatibility with Intel i960 I/O processors and the Intel® 80303 I/O processor.

A robust Intel XScale technology tools environment is now available from Intel and leading third-party vendors to speed the development of intelligent Internet storage, networking, and wireless applications. The comprehensive tools environment features comprehensive software development suites, real-time operating systems, hardware reference platforms, companion chips, debuggers, JTAG and debug tools, logic analyzers, and models. Available tool chains

include compilers, assemblers, and linkers that support specific optimizations for the Intel XScale microarchitecture, the ARM instruction set v.5TE, and Intel® DSP extensions.

Intel also provides the IQ80310 evaluation platform, a development kit that enables hardware and software debugging with JTAG, Ethernet, and serial ports, MICTOR connectors for a logic analyzer connection, and a Macraigor JTAG emulator. The kit includes:

- Cyclone Microsystems* IQ80310 PCI board.
- Core development tools– ARM ADS v1.1* CD-ROM and Cygnus/Red Hat GNUPro CD-ROM.
- Software images for supported real-time operating systems and debug monitors–Wind River Tornado* development environment CD-ROMs, LynuxWorks BlueCat* CD-ROM, and a flash recovery utility CD-ROM.
- Electronic manuals and documentation.

Check the More Info section at the end of this article for links to Intel's Web-based tools resources.

New Directions

Intel is working with the industry to advance a variety of key initiatives that can help deliver improved performance for Internet storage.

- InfiniBand architecture provides improved access to servers through a centralized fabric that can enrich existing and future Fibre Channel SAN deployments. Intel is working with leaders in the enterprise arena to deliver InfiniBand fabrics for Internet Data Center deployment.
- The emerging iSCSI specification defines a protocol to carry storage traffic over existing wide-area Ethernet/IP networks. The encapsulation of block-level data routing over Internet Protocol (IP) enables data to be shared across a variety of storage technologies, ranging from direct-attached SCSI storage devices to Fibre Channel SANs. The Intel® PRO/1000 T IP Storage Adapter connects storage devices located anywhere on an Ethernet network at up to Gigabit speeds.
- The I/O interconnect plays an important role in meeting requirements for performance, scalability, and stability in storage platforms. The multi-drop parallel bus architecture of PCI technology does not provide the scalability needed to sustain the growth of Internet storage over the next decade. Intel and other industry leaders are promoting a third-generation I/O interconnect specification, code-named "Arapahoe," designed to deliver the benefits of full serial architecture for in-the-box expansion. This specification is being defined with wide industry support to meet the emerging and future I/O interconnect requirements of multiple market segments, including Internet storage.

Summary

In the time it takes to read this sentence, thousands of new users will log on to the Internet for the first time, and terabytes of data will flow between storage subsystems. Intel is building on its core competencies in silicon, middleware, and I/O applications to deliver the pre-validated building blocks with multiple levels of integration.

The Intel IOP310 I/O processor chipset, Intel's fourth-generation I/O processor, is an example. The first implementation of Intel XScale microarchitecture, it provides a fully validated solution optimized for intelligent Internet storage. Intel and third-party vendors provide a comprehensive tools environment to accelerate the development of new products based on Intel XScale technology.

In addition to providing silicon, software, and tools, Intel continues to work with the industry to advance new initiatives, including InfiniBand architecture, the iSCSI specification, and the new specification for serial interconnect technology, to help developers deliver breakthrough performance for intelligent Internet storage.

More Info

Check the links below for additional information on the technologies and tools resources mentioned in this article:

- Intel® Internet Storage Building Blocks
- Time to Deliver Internet Storage
- Intel® IQ80310 Evaluation Board
- Intel XScale Microarchitecture
- Embedded Intel® Architecture Communications Reference Designs

Author Bio

Wendy Vittori joined Intel in 1989 as director of marketing and strategic planning for the Supercomputer Systems Division, and later served as general manager for supercomputer systems in Europe. Prior to her present position she was director of operations for Intel's Enterprise Server Group.

Before joining Intel, Wendy held product-line management positions with the Digital Equipment Corporation, was co-founder and CEO of Computers in Medicine (Cambridge, Massachusetts), served as a management consultant to high-technology new ventures, and was a professor of management and entrepreneurship at Northeastern University (Boston, Massachusetts).

Wendy earned an A.B. in philosophy and applied mathematics from Harvard College, an M.B.A. in marketing from Northeastern University, and completed doctoral course work in strategic management and organizational behavior at the Harvard Business School. She has published articles in the field of information management and organization theory.

Software

The Intel® Integrated Performance Primitives API

Stewart Taylor
IPP Chief Architect
Microcomputer Software Lab
Intel Corporation

Overview

There's good news for software developers seeking to increase application performance and decrease engineering costs and time-to-market. That news is the impending release of version 1.1 of the Intel® Integrated Performance Primitives application programming interface (API) and, following it by about a month, the release of version 2.0 beta.

The Intel Integrated Performance Primitives API, also known as IPP, is a low-level software layer, common across server, desktop, and handheld processors and operational in both Windows* and Linux* operating-system environments. An element of the Intel® Performance Libraries, IPP provides developers a broad set of multimedia functionality, including highly targeted functions for mathematics and for signal, audio, image, video, and graphics processing. Moreover, IPP functions are created with highly optimized, processor-specific code that makes it easy for developers to take advantage of processor capabilities that otherwise might be difficult or impossible to access.

Advantages over Earlier Libraries

IPP began as a lower-level version of IPL (Intel® Image Processing Library) and then expanded to include and consolidate primitives from other domain-specific libraries, such as IJL (Intel® JPEG Library) and SPL (Signal Processing Library).

IPP differs from these older libraries in a number of ways. For example, before IPP, when developers needed to use the functions of two or more libraries, they had to install multiple packages. But with IPP there is a single integrated package. More important, with multiple libraries developers had to decipher multiple naming conventions. IPP, in contrast, provides a uniform, easy-to-learn naming convention that assembles function names from four orthogonal components: data type, data size, integer or floating point, and data layout.

Another advantage of using IPP is the efficiency it delivers, through its minimal overhead and its maximum optimization for each of the platforms it supports: Intel® StrongArm* processor, the Intel® Xscale™ processor, the Intel® Itanium™ processor family, and the Intel® Architecture (Pentium®) processor family. For example, a single function call to an older library such as IPL might conceal a number of unexpected conversions or other overhead. By contrast, in IPP all but the most rarely used functions are optimized for a given platform without extra features that can add superfluous code.

Still another advantage is IPP's cross-platform nature. IPP provides similar functionality and a similar or identical calling interface across the platforms it supports. Finally, IPP functions are easy to use in existing applications, because they rely on common operation, require no I/O or local storage, and behave in a manner that is predictable, well-defined, well-documented, and intuitive.

Comprehensive Functionality

The Intel Integrated Performance Primitives API covers a broad range of targeted functionality, including basic and advanced arithmetic functions, small-matrix functions, imaging and video-processing functions, and signal-processing functions, including speech and audio.

Basic arithmetic functions, available in IPP v.1.1, include the following:

- Support: Conj, Copy, Imag, Real, Zero, Set
- Arithmetic: Abs, Add, Convolution, Cross-Correlation, Div, Exp, Ln, Normalize, Mul, Sqr, Sqrt, Sub, Threshold
- Convert: Complex/Real, Integer/Float
- Logical: And, Or, Xor, LShift, Rshift

Advanced arithmetic functions, available in IPP v.2.0, are faster than their single-operation counterparts available in older libraries and faster than the processor operation itself. Highly suitable for scientific algorithms, these functions include the following:

- Basic real elementary: Sin, Cos, Tan, Asin, Acos, Atan, Atan2, Exp, Log, Log10, Pow, Sqrt, Cbrt
- Hyperbolic real: Sinh, Cosh, Tanh, Asinh, Acosh, Atanh
- Secondary real: Expm1, Log1p, Exp2, Exp10, Log2, Frexp, Ilogb, Ldexp, Logb, Hypot, Modf

Small-matrix functions, available in IPP v.2.0, are similar to basic and advanced arithmetic functions but different in their intended application. Whereas the arithmetic functions are primarily for scientific work, the 5,000-plus small-matrix functions in IPP v.2.0 support geometric calculations involving both trigonometry and matrix mathematics and generally targeting graphics work ranging from 3D CAD to games. These functions include the following:

- Vector algebra: Add, Mul, Sub, Scale; Saxpy; CrossProduct, DotProduct; Norm; Linear Combination of Vectors
- Matrix algebra: Add, Mul, Sub, Scale; Gaxpy; Transpose; Det, Inv; Rotate
- Manipulation and conversion
- Solving linear systems: LU decomposition, Cholesky decomposition, single-value decomposition, finding Eigen values, and vectors for matrix
- Least-squares problem: Givens rotations, Householder transformations, QR decomposition

Image-processing functions available in IPP v.1.1, support standard imaging functions as well as computer vision, JPEG, and video compression and decompression. These functions include the following:

- Support: Conj, Copy, Zero, Set
- Arithmetic and logical: Abs, Add, Convolve, Cross-correlation, Div, Exp, Ln, LShift, Normalize, Mul, RShift, Sqr, Sqrt, Sub, Threshold; And, Not, Or; Compare; Phase, Magnitude
- Convert: Pixel/planar, Color, Short to float, Resize
- Filters: User-defined and built-in
- Transforms: FFT, DFT, DCT, Wavelet
- Statistics: Norms, Threshold, Min / Max / std.dev., Mean, PowerSpectr, Moments
- Geometric: Mirror, Rotate, Resize, Remap
- Alpha composite
- Gamma correction
- Image Generation
- JPEG: special color conversions, DCTs, Huffman encode/decode
- MPEG-4: Bitstream processing, motion prediction, inverse quantization
- H.263: Bitstream processing, interblock reconstruction, motion prediction, deblocking filter

Signal processing functions (including speech recognition and audio), available in IPP v.1.1, include the following:

- Support: Conj, Copy, Imag, Real, Zero, Set
- Convert: Polar/Cart, Complex/Real; Integer/Float; Up/Down sample
- Windowing
- Signal generation: random, wave patterns
- Filters: FIR, IIR; Median
- Transforms: FFT, DFT, Goertzel, DCT, Wavelet
- Statistics: Norms, Threshold, Min/Max/std.dev., Mean, PowerSpectr
- Audio: A-law/mu-law, preemphasize
- Speech recognition: Two-dimensional arithmetic and vector functions, basic recognition functions, Gaussian mixture functions

Summary

The Intel Integrated Performance Primitives API (IPP) offers multimedia developers a faster and more efficient way to incorporate functions for mathematics and for signal, audio, image, video, and graphics processing into their products. With the release of IPP v.1.1, to be followed shortly by the first beta release of IPP v2.0, developers can take advantage of a set of primitive functions that are cross-platform and optimized for the Intel StrongArm processor, the Intel XScale processor, the Intel Itanium processor family, and the Intel Architecture (Pentium) processor family. Designed to simplify the accessing of processor capabilities and to significantly reduce overhead, IPP can help developers bring their products to market faster and at a lower cost.

More Info

Developers interested in the capabilities of IPP can visit the IPP Web site to learn more about the technology, to download a free evaluation copy of IPP v.1.1, to access more comprehensive information from an Intel SRM or BDM, or to apply for the IPP 2.0 beta. Note also that IPP v.1.1 or v.2.0 is packaged for desktops in two ways: (1) one DLL/processor and a separate dispatcher, and (2) one static library with a dispatcher inside. IPP also provides static libraries for handheld or other wireless devices and in whatever form is both cost- and royalty-free.

Author Bio

Stewart Taylor is the IPP chief architect in the Microcomputer Software Lab at Intel Corporation. Over the past few years he has been designing libraries for Intel, including the Intel Image Processing Library. He has been at Intel since 1992, when he was the developer of the original video codec for the ProShare® conferencing software. Stewart has a B.S.E. from Princeton University, and holds an M.S. and is a candidate for a Ph.D. in the E.E. department at Stanford University.

Automatic Debugging of Multithreaded Programs

Henry A. Gabb
Staff Parallel Applications Engineer
Microcomputer Software Lab
Intel Americas, Inc.

Overview

One of NASA's most successful missions, the Mars Pathfinder, was nearly crippled by a concurrent programming error. Were the programmers at the Jet Propulsion Laboratory (JPL) careless? Hardly. JPL has some of the world's best programmers. The problem is that it is more difficult to debug and test multithreaded programs than it is to debug sequential programs. In a multithreaded program, independent tasks are mapped to threads, which execute concurrently. And, thread scheduling is typically handled by the operating system, not the programmer.

In spite of the added complexity, the performance requirements of today's applications (such as games, video and audio encoding, voice recognition, etc.) are forcing many developers to use multithreading. Developers are now using multithreading not only to express concurrency but also to take advantage of parallel processing. However, the temporal component of multithreaded programming is new to many programmers and can be confusing. This means that more bugs make it into the end product. Programming assumptions (such as execution order) that would be correct in a sequential program lead to errors in a multithreaded program. Thread-related errors are also more difficult to consistently reproduce, so debugging time is magnified.

Thread-aware debuggers have been available for years. However, manual debuggers require that the developer examine all thread states in order to find errors–assuming of course, that the error can be reproduced in the debugger. In contrast, the Assure for Threads tool automatically finds multithreaded programming errors, such as deadlock and data races. The tool helps programmers test and debug existing multithreaded programs. It also helps programmers who are developing new multithreaded applications.

Concurrency

Thread libraries are no more difficult to use than other libraries. However, thread libraries are designed to allow independent tasks to execute concurrently, but even independent tasks must sometimes modify shared data.

Synchronizing access to shared data is the programmer's responsibility. In the absence of explicit synchronization, the operating system is free to schedule threads in any order, no matter how inconvenient this scheduling is to the programmer or the end user. But as programmers, we expect statements to execute in the order in which they are coded. This assumption is the primary cause of errors in a multithreaded program. When multiple threads call the same function simultaneously, statement execution can be interleaved. For example, if the function modifies a static variable, a data race is possible, which can corrupt the data and produce incorrect results.

It is the asynchronous nature of concurrent programming that makes debugging more difficult. If the developer is lucky, errors related to concurrency cause program failure. If the developer is not lucky, these errors can cause occasional failures that are difficult to reproduce. They can also cause subtle numerical errors that are difficult to find. So, in the developer's environment, such errors may be extremely rare and can go unnoticed. When the program is run in the end-user's environment, the error may occur more frequently.

Assure for Threads

Assure for Threads is the first tool of its kind for automatically debugging multithreaded programs. Assure for Threads can detect subtle data dependencies, memory leaks and illegal memory accesses, and deadlocks caused by incorrect locking hierarchies, lost signals, and dangling locks. One of the main benefits to using this tool is that it requires no recoding. Programmers are not required to manually instrument the code prior to analysis. Instead, Assure for Threads automatically instruments the code at compile-time. Another major feature is that the tool pinpoints and highlights the exact lines of code responsible for errors, even when the source lines causing the error are in different parts of the program. The tool supports both C and C++ programs.

Fall 2001 Intel Developer Forum

Developers saw a demonstration of Assure for Threads at the Fall 2001 Intel Developer Forum (IDF) Conference. The conference included both a lecture and a lab for multithreaded programming and debugging. The lecture covered basic library functions for creating and synchronizing threads. During the session, Assure for Threads was used to debug seemingly correct programs that contain subtle errors when executed with more than one thread. In the lab, developers had the opportunity to use Assure for Threads to find the errors in seemingly correct code.

Parallel Applications Center

Intel's Parallel Applications Center is available to Independent Software Vendors (ISVs) for confidential test and analysis of multithreaded programs. By testing released products, developers often find hidden and sporadic problems that are currently affecting end users. The Parallel Applications Center also helps ISVs thread programs to boost performance. Developers should contact the Intel® Parallel Applications Center to find out about services that are now available to ISVs who want to take advantage of multithreading on Intel® parallel architectures.

Summary

No matter how fast processors become, they will never be fast enough. The computational requirements of today's games, for example, push the newest high-speed CPUs to the limit. Some digital video encoding algorithms already exceed the capability of individual processors. Multithreading lets developers apply more processors to a computation to reduce processing time and dramatically boost performance.

Concurrent or multithreaded programming has been around for a long time, but it has always been more difficult than standard, sequential programming. Because of the added complexity of creating or updating a multithreaded program, debug time can be extensive. While processor time is cheap, programmer time is not. Developers should let Assure for Threads automatically debug their threaded programs. Product managers should also add Assure for Threads to their quality assurance and product-release criteria.

More Info

Developers should visit the Assure for Threads Web site for more information about the debugging tool. The Web site includes descriptions of the tool’s interface and debugging features, as well as a tutorial with sample multithreaded programs. The site also offers information about developer services available through the Intel Parallel Applications Center.

Author Bio

Henry A. Gabb has been with Intel one year as staff parallel applications engineer. He works in Intel’s Parallel Applications Center, which helps ISVs take advantage of Intel's parallel architectures. Prior to joining Intel, Henry was director of scientific computing at the U.S. Army Engineer Research and Development Center, a Department of Defense High-Performance Computing facility. He also worked in the Biomolecular Modeling Laboratory at the Imperial Cancer Research Fund (London, England), where he developed software for computer-aided drug design. Henry did his postdoctoral research at l’Institut de Biologie et Physico-Chimique (Paris, France), where he developed computer models to simulate conformational changes in biomolecules.

Henry has been awarded research fellowships from Le Ministère des Affaires étrangères in France, the Imperial Cancer Research Fund in the United Kingdom, and the European Molecular Biology Organization. He has published over 20 peer-reviewed publications, as well as five invited-review articles. Henry received his B.S. in biochemistry from Louisiana State University. He received his Ph.D. in biochemistry and molecular genetics from the University of Alabama, Schools of Medicine and Dentistry.

Web Development

Intel Drives Richer Web Services through e-Business Standards

Joel Munter
Senior Software Engineer
e-Business Solutions Lab
Intel Corporation

Christian Thomas
Senior Software Engineer
e-Business Solutions Lab
Intel Corporation

Overview

For today's Web service registries, the de facto specification is Universal Description, Discovery, and Integration (UDDI). One successful and increasingly popular UDDI implementation already exists in the Web. By August of 2001, the UDDI Business Registry included over 5,200 businesses and 7,800 services. More businesses and services are registering every day.

Foreseeing the need for a UDDI Business Registry, Intel previously created the Dynamic Service Discovery (DSD) architecture. In addition to the basic UDDI registry functionality, the DSD architecture includes client subscription, rich query, leasing, and binary proxies—features not yet available with the current version of UDDI. The unique capabilities of DSD have now been isolated and layered onto the UDDI Business Registry.

The UDDI Version 2 registry shows that service discovery can be practical, cost-effective, easy to use, and a valuable way to promote Web services. Work is now underway for UDDI Version 3, to make e-Business registries more accurate and more effective. This work includes plans to address improved query facilities, caching and subscriptions, security, and internationalization. In particular, Intel is advocating two important enhancements to UDDI: rich query capability (via comparison engine) and client subscription. For future versions of UDDI, Intel is advocating other enhancements, such as communications proxies, client-service leasing, and peer-to-peer discovery. Intel is also coordinating with RosettaNet to register some of their key Partner Interface Processes into the UDDI Business Registry. Once the UDDI specification includes DSD and similar features, the specification will be even more attractive and functional for developers and users.

What is UDDI?

UDDI is a service-discovery specification that allows registered businesses and Web services to be discovered by any client using the defined UDDI API's. Businesses register information about themselves and their services in the UDDI Business Registry, which is like a rich, global yellow-page directory. This directory or registry contains enough information about each registered business and service to let users begin interacting directly with those entities.

The UDDI Business Registry is a true registry. Actual Web-service specifications or supporting documents are not stored within the registry. The registry acts only as an index for Web services. The interaction between a user and the Web services is independent of UDDI.

Businesses and services can also be registered in the UDDI Business Registry by third-party registrars or brokers. These registrars support the standard UDDI API functions and may offer services not yet supported by the UDDI specification.

Once a business or service is discovered, the user can obtain enough information about the business to begin interacting directly with that organization. For example, with UDDI, a broker can resolve technical implementation issues with other brokers, so clients do not have to worry about such issues. UDDI also resolves communication issues between businesses. Without UDDI, a client using Salutation or UPnP may not discover an E-speak service, and vice versa. However, with UDDI, businesses and services listed in the UDDI Business Registry can be discovered, regardless of the application used.

Global Registry

UDDI links individual registries into a single, global, public UDDI Business Registry. Currently, two registries have been linked: the IBM registry and the Microsoft registry. Hewlett-Packard also expects to add its registry to the UDDI Business Registry soon. Each individual registry owns a subset of the globally available data, but that data is shared with or replicated in the other registries in the UDDI Business Registry. A client or user interacting with, for example, the IBM registry, will actually have access to the entire data set of the UDDI Business Registry.

Business and service discovery and information drilldown type interactions between a UDDI registry and another Web service or client are then performed via the Simple Object Access Protocol (SOAP).

The UDDI specification can be used for private registries, not just for registries listed in the global, public UDDI Business Registry. Such private registries are useful in vertical markets, such as a market for only auto parts or for only health care. Private registries are also useful within a large organization or company, such as Intel. Another example of a private registry would be a group of e-Business partners that want to share information and Web services, but do not want this information available to the general public. Private registries can be users or clients of the UDDI Business Registry.

Subscription

As a lead member of the 15-company working group, Intel is helping define how client subscription would work in UDDI. Client subscription lets a client register interest in a particular business or service. For example, a client might want to know when a new business or vendor becomes registered that provides wholesale bicycle sprockets.

Subscription lets the client get real-time information even as data changes within a registry or with the UDDI Business Registry. The alternative is to have a client continually polling a registry with a particular query, adding an unnecessary workload to the registry and potentially inundating the network. With UDDI and support for subscription, the client can register a one-time or standing request for the type of business or Web service that it is interested in interacting with (such as that bicycle sprocket vendor).

Rich Query

Intel is also helping define how rich query would work in UDDI. Rich query capability would let clients discover information and services in a more robust manner. For example, a client might want to discover the nearest 1,200-dpi printer. Whether the client uses the number "1,200" or spells out "twelve hundred" shouldn't matter. The query should be successfully passed and understood with either term.

One way to perform rich queries is to use comparison functions. Comparison functions can live in the Web service or can be downloaded into the actual registry, such as the UDDI Business Registry. Note that the UDDI Business Registry does not need to understand context, such as what a 1,200-dpi printer is or can do. Instead, the UDDI Business Registry can rely on the content expert to offer the client the context, description, specifications, etc., directly once the client and business connect.

Another example of a rich query is a request to find the nearest gas station within one mile of your current GPS coordinates. You wouldn't expect the overall UDDI to understand all the issues involved in refilling your fuel tank. Instead, UDDI lets you use a rich query function to discover a list of vendors that have made that information available to you.

Other Capabilities

The UDDI specification provides a well-defined base for discovery of businesses and services. As each new version of UDDI is defined, the functionality will continue to evolve to include the discovery features that businesses and clients demand.

Until these features are available from UDDI, a registrar or broker can add that functionality for clients. For example, UDDI Version 2 does not have subscription capabilities, and has only limited query capabilities. Another way for companies to enhance the basic functionality of UDDI is to incorporate Intel's DSD complementary architecture (which includes rich query) and other features into a registrar-type client. Enhancements currently being proposed for UDDI Version 3 include client subscription (subscription notification) and rich query capability.

More Info

The UDDI specification, recommended practices, FAQs, and so on are available online at the UDDI Web site. The DSD architecture, which complements UDDI, is described in the Intel white paper titled “Dynamic Service Discovery Executive White Paper (UDDI Enhancement Proposal).“ The SOAP standard, which is used to control interactions between UDDI clients and services, is also available online.

Summary

Businesses are using automated services to do more and more of their interacting. UDDI makes it easier for users to automatically discover other businesses and services. In fact, UDDI is easy to use and implement. Developers should look at incorporating DSD features into their Web services. Intel's DSD architecture adds significant value to UDDI Web-service interactions. While Intel is proposing that some of these capabilities be included in UDDI Version 3, most features are likely to be included in later versions of the UDDI specification.

Author Bio

Joel Munter has been with Intel for over eight years as a senior software engineer. Currently, he is the lead Intel representative to UDDI from the Intel e-Business Solutions Lab. Previously; he worked as lead developer of ENGIS System, which was used to manage Certificate of Conformance/Quality data and general lot-quality issues. He has also been CS/IS lead developer and has been project manager for the Intel Environmental Health and Safety Workstation effort. Joel has received several awards, including three Division Recognition Awards for his work on the EHS Workstation project, EDI implementation, and VAX-to-NCR transition project. He also received the Best Paper award at the Intel Software Development Conference. Joel received his B.S.E. from Arizona State University, with a combined program of mechanical and aerospace engineering.

Christian Thomas has been with Intel over a year as senior software engineer and Web service technologist in the eBusiness Solutions Lab. Previously, Christian worked as lead architect for Intel's Dynamic Service Discovery specification. Along with a co-author, Christian holds two patents and has filed three others. He has also received the Technical Achievement Award from Honeywell for the design of the Honeywell Communications Infrastructure (HCI) product. Christian received a B.S. in computer science from Union College.

Wireless

Automated Client-Based Layer-3 Switching

Nikhil Deshpande
Applications Engineering Manager
Intel Labs
Intel Corporation

Overview

Continuous, transparent, mobile connections—that's the ideal for today's corporate networked communications. Mobility, in particular, is key. Laptops moving from desk to conference rooms to car must continue to support wired LAN while also supporting wireless IEEE 802.11b LAN for mobile connections.

Continuous and transparent connections can be maintained in layer 2 of the network stack. However, this connectivity is continuous and transparent only within its own subnet. Continuous connections at layer 3 and higher are difficult to maintain. For example, at layer 3, mobile users may be trying to connect to access points on a variety of subnets.

Intel has developed a new, client-based approach for wireless connectivity called Adapter Switching (AS) software. AS software deals with wireless connectivity across multiple networks, such as wired LAN, IEEE 802.11b LAN. The AS software also supports several current and legacy operating systems such as Microsoft Windows* 98, Windows 98 SE, Windows Me, and Windows 2000.

Client-Based Mobility

In the future, the issue of mobility will likely be solved by mobile IP software. Mobile IP lets users move from place to place and still maintain a connection. However, mobile IP technologies also require that additional infrastructure be placed in the network.

Today’s AS software is a client-based solution that does not require any additional infrastructure in the network. This software lets users continue a connection without rebooting the system. The connection may not be continuous or unbroken, but users can re-establish the connection simply by refreshing or restarting the network application, rather than by rebooting the system. For example, Microsoft Internet Explorer includes a refresh feature that can be used to re-establish the connection without closing and restarting the application. This means that users can, for example, move a laptop from a wireless subnet at a conference to a wired desktop environment, make the hot physical network connections for the wired LAN, and refresh the network application to continue work without interruption. The AS software automatically switches the network interface card (NIC) as appropriate for the available access points.

Preferences

In most of today’s laptops, users cannot specify a preference for wired or wireless connections when both types are available. However, wired connections can have much higher the bandwidth that wireless connections have. When a wired LAN is available, users should be able to specify that the wired LAN be used before a wireless connection is established.

The AS software lets users specify a preference for wired or wireless communications when both types of LANs are available. User preferences are then tracked by the AS policy manager. When an NIC is disconnected and/or reconnected to a network or subnet, the software uses various components (such as the link monitor and policy manager) to determine which type of connection should be made.

Automated Switching

Intel's AS software is, essentially, automated layer-3 switching software. When an event occurs (when a media card becomes active), the AS software becomes active. The link monitor tracks available links and the IP addresses of those links; the policy manager changes the routing table if necessary; and so on.

When a new interface becomes available (for example, an NIC reconnects), the operating system (OS) automatically assigns it an interface IP address. When a data packet is initiated, the OS compares the destination IP address to the interface IP address and subnet mask. If the destination IP address matches the interface IP address and subnet mask, the OS selects that particular interface for communication.

If there is no subnet match, the OS uses secondary criteria to select an interface for communication. The secondary criterion is the metric value of each network interface card (NIC). At this point, the link monitor checks with the policy manager to see whether user preferences were specified for the installed NICs. The policy manager then decides whether any changes should be made in the routing table. If necessary, the metric value of the lower priority card is changed to a higher number (such as from 1 to 2). The OS then routes the packet on the interface with the lowest metric value.

You might ask what happens to data traffic that was previously going through one NIC when another NIC becomes active. With AS software, the previous data traffic continues to flow through the previously active NIC. New traffic is routed through the NIC with the lowest metric value. A side benefit of this technique is that users can experience increased bandwidth by sending data through both wired and wireless pipes at the same time.

Metrics

An NIC's metric value isn't changed in the routing table until and unless there is a need to do so. For example, a laptop may have only one NIC installed. In this case, that NIC's metric value is not changed because that card is the only one that can be used, regardless of user preferences. In the same vein, if two cards have default metric values that already match the user preferences, those metric values also remain unchanged in the routing table. If all metric values are identical, the NIC is picked at random by the OS.

Intel's AS software is designed for Intel® NICs. The metric value of non-Intel NICs is set by the operating system. In contrast, the metric value of Intel NICs can be changed by the AS software as specified by user preferences. Developers who use Intel NICs will find that the AS software extends the feature set of their laptop networking products.

Summary

Intel's Adapter Switching software provides automated layer-3 switching. AS software monitors the network link status on two or more NICs, lets users specify LAN preferences, supports several legacy systems, and automates network switching. Intel's AS software:

- Lets users define preferences for wired and wireless connections when multiple choices become available
- Enables preference-based switching on legacy operating systems such as Windows 98, Windows 98SE, Windows Me, and Windows 2000
- Automates IP address assignment when a card is inserted into a computer running Windows 98
- Lets users switch from one access point to another on different subnets

AS software is a client-based solution for laptops that may act as a foundation for future seamless mobile communications. Intel's AS software solution can be easily extended to hand-held devices, such as PDAs (personal digital assistants) and pocket PCs. Other, more futuristic, server-based solutions will offer even more seamless roaming connectivity.

More Info

Several white papers on wireless LAN and information on wireless connectivity and related topics are located at the Intel Wireless Internet Solutions site.

Author Bio

Nikhil Deshpande is applications engineering manager in the Mobile Data Service area of Intel Architecture Labs. Although he has been with Intel for just a year, Nikhil has worked in wireless technologies for six years. Prior to Intel, he developed algorithms for decoding CDMA/TDMA/GSM signals and worked on 3G wireless technologies such as W-CDMA, CDMA2000, and EDGE. He researched CDMA forward-link signal generation. Currently, Nikhil is working on strategic planning and next-generation collaboration applications for mobile professionals. He has taught seminars and college courses in wireless technology, has published numerous papers, and has sat on standards bodies such as 3GPP and 3GPP2. Nikhil has 10 patents pending in digital communications and wireless field. He received a B.S. in electronics and telecommunications engineering in India. Nikhil received his M.S. and Ph.D. in electrical and computer engineering from Portland State University.

—End of Intel Developer Update Magazine Issue 24—